

Part IV

Real Estate Development

Part IV will focus on the financial analysis of real estate development, beginning with the closely related but more fundamental topic of land valuation, including the real options model.

The real estate development industry, along with the space market and the asset market, is the third major element in the real estate system modeled in the Quadrant Model of Chapter 6. Real estate development is a large and important sector of the US economy, with construction generally accounting for over six percent of the GDP and non-residential construction alone employing over 800,000 workers. More importantly, real estate development tremendously affects the physical and social environment. It is how cities are built. It is important to understand the industry from an economic and investment perspective for reasons that go beyond just economics and investment.

The real estate development industry is the engine of entrepreneurial activity that assembles and applies the financial, physical, and legal/political resources to construct newly built space. Development is a complex and creative function that, at its best, displays great vision and, at its worst, enormous greed but, in almost all cases, considerable risk-taking on the part of developers and/or their financial backers. Within the development-industry component of the real estate system, a crucial comparison is made. Current development costs, including construction and land costs (incorporating necessary profit for the developer), are compared against current asset values. If asset values equal or exceed development costs by a sufficient margin reflecting timing risks, then development will proceed, thereby adding to the physical stock on the supply side of the space market. Of course, a key component of development costs is the opportunity value of the land. This opportunity value is determined in the real estate asset market, as land is a type of real estate asset. Of course, the asset market values land as potentially developable sites, more formally as call options on the developable building(s), which in turn are related fundamentally to the space market that each site could serve once developed. And development takes time, which requires that the development industry be forward-looking. The developer succeeds only if the newly developed property's value exceeds its total development cost at the time of completion of the project, which may take several years.

Real estate development is, in many respects, the most exciting and dynamic aspect of real estate investments. In Part IV, we will take you on a sophisticated tour of the underlying framework that governs development decisions and values from a rigorous economic perspective, consistent with the foundations and principles taught in the earlier parts of this book.



Taylor & Francis
Taylor & Francis Group
<http://taylorandfrancis.com>

16 Real Options and Land Value

CHAPTER OUTLINE

- 16.1 The Call Option Model of Land Value
- 16.2 Simple Numerical Example of OVT Applied to Land Valuation and the Development Timing Decision
- *16.3 A Rigorous Model of Option Value
 - *16.3.1 The Arbitrage Perspective on the Option Value
 - *16.3.2 The Certainty-Equivalence DCF Perspective on the Option Value
 - *16.3.3 How the Option Valuation Model Works
- 16.4 The Binomial Model of Option Value
- *16.5 A Perpetual Model in Continuous Time
 - *16.5.1 The Classical Samuelson-McKean Formula
 - *16.5.2 General Implications of the Model for Development Timing and Land Speculation Investment Risk and Return Expectations
- 16.6 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- Completed development value must exceed costs by a sufficient margin to induce development. Until then there remains an option on land.
- The call option and how this concept can be used to understand the value of land and helps us understand why it is sometimes more beneficial to delay development, even though feasible development is currently possible.
- The binomial option valuation model and how and why it works.
- The Samuelson-McKean formula and how it can be used to shed light on land value, development timing, and the opportunity cost of capital for investment in land speculation.
- Some of the insights that option valuation theory provides for understanding real estate development behavior.

Land value is probably the most fundamental topic in all of real estate in general and real estate development in specific. In our discussion of the real estate system in Chapter 6 of Part II, we saw land value's pivotal role as the key link between the asset and space markets in the real estate development process. In addition, we saw the crucial role of land value in urban economic analysis and the shape and dynamics of real-world urban development. Indeed, land is the fundamental defining characteristic of real estate, and the nature of land valuation helps to define the investment characteristics of most major real estate assets. In this chapter, we will introduce you to a very useful (and Nobel Prize-winning) tool developed during the last few decades to help understand, analyze, and evaluate land and the real estate development process. This tool is **option valuation theory (OVT)**, especially the branch of that theory known as **real options**.

At the same time, it should be acknowledged upfront that option value theory can be complicated and has few direct uses in real estate development. Thus, be warned, this chapter is mostly for students who want to understand land value and its implications for urban form on a very fundamental and rigorously formal level. This can give you deep insight and appreciation that may be useful, certainly in the academic world, and in some career paths in the investment industry. The lessons in this chapter will be of interest not only for their substantive implications about land value but also for the methodology and analytical tools that will be introduced here (which are more generally applicable beyond just land valuation to subjects such as financial options and derivatives). If you are simply interested in developing an office tower, you have our blessing to skip directly to Chapters 17 and 18. You should be able to follow those chapters without going through this one first.

16.1 THE CALL OPTION MODEL OF LAND VALUE

In finance, an **option** is defined as follows:

An option is a right without obligation to obtain something of value upon the payment or give up something else of value.

The person having such a right is referred to as the owner or holder of the option. The asset obtained by **the exercise** of the option is often referred to as **the underlying asset**, while that which is given up is referred to as the **exercise price** of the option. The option holder has the right to decide whether or not to exercise the option. If the option can be exercised at any time (before an expiration date or maturity, if the option has such), then the option is said to be an **American option**. If the option can only be exercised on its expiration date (not before), it is considered a **European option**. In any case, option exercise is irreversible because, in the act of exercising the option, the option itself is thereby given up; that is, an option can only be exercised once.

OVT consists of a body of theory and methodology for quantitatively evaluating options. Included as an integral part of such valuation is the problem of specifying the conditions when it is optimal to exercise the option if it is an American option. The classical types of options in financial economics for which OVT was first developed are the stock options and warrants traded in several public exchanges and widely issued to corporate executives as an incentive component of their compensation. For example, the classical common stock **call option** gives the holder the right, without obligation, to purchase at a stated price per share a specified number of shares of the common stock of a specified company on or before a certain date.

The term “*real options*” refers to the study of options whose underlying assets (that is, either what is obtained or what is given up on the exercise of the option) are real assets (i.e., physical capital) as opposed to purely financial assets. For example, a building or a factory is a real asset, whereas shares of common stock or a release from a mortgage debt obligation are financial assets. The default option in mortgage debt (the mortgagor’s put option) is an example of a real option because the borrower is giving up real property. However, a deeper and more fundamental level of application of real option theory to real estate is to apply the option model to the land itself. When we do this, we go to the very heart of the real estate system because OVT can now shed light directly on the relation between land value and the timing and nature of the development of buildings on the land.

Applying real options theory to real estate may be termed the **call option model of land value**. In this model, land is viewed as obtaining its value through the option it gives its owner to develop a structure on the land. The landowner can obtain a valuable rent-paying asset upon paying the construction cost necessary to build the structure. More broadly, the landowner’s option also includes the option to demolish and/or redevelop any existing structures on the land. However, unless the existing building is quite old or small, or the development of the city and neighborhood has rendered the existing structure inappropriate for the best use of the location, the cost of demolishing the existing building (in particular, the *opportunity cost* of the forgone revenues that building could earn) will normally be so great as to minimize the redevelopment option value. Thus, the exercise of the option is essentially irreversible; the option is given up through its exercise. The option model

of land value is therefore seen to be most applicable either to vacant (or nearly vacant) land, such as surface parking lots, or else to land in transition zones where the highest and best urban use of the land is changing.¹ In any case, it is important to remember that the real option that is viewed as giving the land its value is, essentially, the *land development option*.

The call option model of land value can have such a fundamental and central place in furthering our understanding of real estate because it relates directly to some crucial links in the big picture of the real estate system, as we described in Chapters 6 and 7. Land is the characteristic component distinguishing real estate from other types of capital assets.

The call option model allows us to better understand and quantify land value. Real option theory highlights and formalizes in a rigorous way the important link between land value and real estate development. If you recall from Chapter 6, our depiction of the big picture of the real estate system depicted in the Quadrant Model, you will remember the crucial role of the development industry in providing the long-run link between the real estate asset market and the real estate space usage market. The lower-left quadrant in that model links the capital asset market to the actual supply of built space. From Chapters 6 and 7, you will also recall the crucial role of the land market in determining urban spatial form. Finally, it is difficult to apply the traditional discount rate approach to DCF analysis to evaluate development projects because it is difficult to know what the appropriate opportunity cost of capital is to use as the discount rate.² These issues can be addressed rigorously and quantitatively and integrated as never before, using real options theory and the call option model of land value.

16.2 SIMPLE NUMERICAL EXAMPLE OF OVT APPLIED TO LAND VALUATION AND THE DEVELOPMENT TIMING DECISION

To gain a more concrete understanding of OVT's fundamental application to real estate, let's walk through a very simple example of how the development option affects land value and the development timing decision. Consider a vacant land parcel and its potential development as summarized in Exhibit 16-1. Suppose a building worth \$100 million today can be developed on the land at a construction cost of \$88.24 million for an immediate profit of \$11.76 million.³ Suppose the development option does not expire today but will exist for another year. Further, suppose that it is reasonable to expect that, if we wait until next year, there is a 70 percent chance that the market will improve and a similar building, newly completed next year, would, in that case, be worth \$113.21 million. On the other hand, there is a 30 percent chance that the market will decline, resulting in a value of only \$78.62 million for the new building next year.⁴ In either case, construction costs will have increased by 2 percent to \$90 million next year.

¹ Recall our neighborhood and property life cycle discussion in sections Chapter 7.

² Section 16.3.2 below will solve this problem formally and practically with a more general method of performing DCF valuation.

³ For now, we will ignore the time that it takes to build; that is, we will assume construction is instantaneous, such that by paying \$88.24 million of construction cost today, we could immediately have a complete, fully operating building worth \$100 million. Note that the present value of this building today includes the present value of the net rental income the building would be expected to provide next year (as well as all subsequent years). It does not, however, include the value of any rent received during the current (just past) year. (In other words we are working with "ex-dividend" asset values). In practice, since the building we could develop on the land does not yet exist, we can only observe the \$100 million current value of such a building by observing the current value of similar buildings that do exist and by extrapolating their rents occupancy, operating expenses, and capitalization rates to our subject development (perhaps with suitable modification based on our understanding of the relevant real estate market). In effect, we are assuming the existence and observable value of a "twin asset" to the one we would build. This type of "twin asset" assumption actually underlies all DCF analysis, as the OCC used in the discount rate is obtained by observing required returns for "similar investments" (i.e., effectively, "twin assets").

⁴ These numbers may seem curious. In fact, they represent a plausible scenario consistent with the value of a newly completed building being worth \$100 million today. Note that the expected value of a newly completed building next year is \$102.83 million $[0.7(\$113.21) + 0.3(\$78.62)]$. This implies an expected annual growth rate of 2.83% in the value of a new building, roughly consistent with typical inflation expectations in the early 2000s. Suppose a new building completed today would generate net rental income next year equal to 6% of its ex-dividend value at that time. Then, the building would generate a net rent of $6\% \times \$102.83 = \6.17 million next year. Thus, an investor buying the newly completed building today would obtain a total expected value next year of $\$102.83 + \$6.17 = \$109$ million. A newly completed building value today of \$100 million would then be equivalent in our scenario to assuming that 9% is an appropriate opportunity cost of capital (OCC, or risk-adjusted discount rate for DCF valuation) for investment in completed buildings and that 6% is a plausible

EXHIBIT 16-1

Numerical Example of Option Premium Value Due to Future Uncertainty in Built Property Value

(Values in \$ millions)	Today	Next Year
Probability	100%*	30% 70%
Value of developed property	\$100.00**	\$78.62 \$113.21
Construction cost (excludes land cost)	\$88.24	\$90.00 \$90.00
NPV of exercise	\$11.76	\$11.38 \$23.21
Future values (actions)		\$0 \$23.21 (Don't build) (Build)
Expected value of built property (Probability x Outcome)	\$100.00** (1.0 x 100)	\$102.83 (0.3 x \$78.62 + 0.7 x \$113.21)
Expected value of option (Probability x Outcome)	\$11.76 (1.0 x 11.76)	\$16.25 (0.3 x \$0 + 0.7 x \$23.21)
PV(today) of alternatives @ 20% discount rate	\$11.76	\$13.54 = \$16.25/1.20
Land Value Today = MAX(11.76, 13.54) = \$13.54		
Option Premium = \$13.54 - \$11.76 = \$1.78		

* Today's values are known for certain (100% probability) because they can be directly empirically observed.

** Today's value of \$100 million is for a building already completed today and includes the present value of the building's expected net rental income next year.

Assuming a built property OCC of 9%, the present value today of a forward claim on a building to be completed next year would be $E[V_1]/(1 + OCC) = \$102.83/1.09 = \94.34 million.

Now, we can see how the flexibility allowed by the lack of obligatory exercise of the development option enables the landowner to take advantage of future possibilities. The landowner can simply choose not to develop if the downside \$78.62 million value outcome occurs next year. (The NPV of development at that point would be $\$78.62 - \$90.00 = -11.38 < 0$, while the NPV of doing nothing would be zero.) Yet the possibility of developing under the upside \$113.21 million contingency would provide a net profit of \$23.21 million at that time (the upside building value minus the \$90 million construction cost). This gives an expected value for the development option next year equal to \$16.25 million, even though there is only a 70 percent chance that the favorable scenario will occur $[\$16.25 = (0.7)(\$113.21 - \$90.00) + (0.3)(\$0)]$. Thus, the option enables the landowner to avoid much of the negative consequences of the downside outcome of future market volatility while still retaining the ability to profit from the upside.

Now, consider the expected NPV of the development project as of the time of development under the two mutually exclusive alternatives of building today or waiting and building in 1 year. The NPV in the first case is the \$11.76 million we noted earlier. The NPV in the second case would be the \$16.25 million we just calculated.

But this second NPV is for one year in the future. Furthermore, it is an *expected* NPV based on our expectations today about the situation, which is one year in the future. As there is uncertainty about what the future will bring, this future NPV is risky, whereas the NPV of immediate development today is known for certain. We must account for the time and risk differences between the two NPVs to compare the two alternatives. This is traditionally done by applying a risk-adjusted discount rate to the second NPV in order to discount it back to a present certainty-equivalent value comparable to the first NPV. The appropriate discount rate is the opportunity cost of capital (OCC, or the investment market's expected total return) of the speculative land investment. As such, land investment is generally considered quite risky, and investors would probably require a rather high expected return, say, 20 percent per annum. This would give us a present value of \$13.54 million for the second NPV, calculated as $\$16,250,000/1.20$.

TEXT BOX 16.1 THE NPV RULE AND (RE)DEVELOPMENT OPTION VALUE

One of the first points that is often highlighted from real options theory is its implications regarding the basic NPV rule we introduced in Chapter 4. It is often said in the corporate capital budgeting literature that option theory shows that the classical NPV rule of investment is too simple. Option theory suggests that it will usually not make sense to invest in a project with a small positive net present value (NPV), as this is quantified in the typical corporate capital budgeting application. Instead, it will usually make more sense in the corporate context to wait until the NPV is substantially positive before investing.

However, a careful reading of the NPV investment rule suggests that option theory does not really negate or modify the old rule, as we have stated in this text. We simply have to be careful about how we apply the rule. Recall from Chapter 4 that the NPV rule is to invest so as to *maximize* your NPV. We don't just invest in any and all positive NPV projects when there are mutually exclusive alternatives. Instead, we pick the alternative that has the maximum NPV. The irreversibility of construction projects means that investing today excludes the possibility of investing later (e.g., in the same or another construction project on the same land). Building today versus building next year are mutually exclusive alternatives for a given site. Maximizing the NPV would require selecting the construction timing that has the highest NPV, which is discounted to today. Applied this way, the NPV rule should still hold.

Another way of seeing this point that is more relevant for real estate applications is to include the current value of the option premium as part of the cost one is incurring by exercising the development option. Whenever one considers the costs and benefits of an investment, one must be careful to include *all costs* and benefits. For an irreversible project, the costs include the option premium given up by undertaking the project. The value of this option premium may often take much work to see and quantify in the case of industrial corporate capital projects, in which highly unique and proprietary types of equipment or physical assets are to be built. However, in the case of the typical real estate development project, the option premium should be more readily observable (albeit still not perfectly so) in the current market value of the land. In a well-functioning land market, the development option premium should normally be included in the current market value of the land. This land value is always an *opportunity cost* of the development project (even if the developer already owns the land).

Thus, on a present certainty-equivalent basis, the NPV of the first alternative (to build today) is \$11.76 million, while the NPV of the second alternative (to wait and possibly build 1 year from today) is \$13.54 million. As these two values are now directly comparable and mutually exclusive, the value-maximizing decision is clearly to wait and develop next year. As the landowner has the right to make this decision, this implies that the land value today must be \$13.54 million, not the \$11.76 million current exercise value.⁵ The difference, \$1.78 million, is what is called the “**option premium**,” which clearly derives from the flexibility provided by land ownership to the owner to develop the land at whatever time she chooses, reflecting the land’s option value.

Note that if the land is worth \$13.54 million, then the NPV of the best current development project, *including the opportunity cost of the land*, is negative (equal to negative \$1.78 million, which is the \$100 million current newly built property value less the \$88.24 million construction cost,

payout rate (or 6.17% a plausible forward-looking “cap rate”) for such buildings. Such rates were indeed typical of the office building investment market in the United States in the early 2000s.)

⁵ In a competitive land market, the price of the land would presumably be bid up to \$13.54, as any landowner would have the right to postpone development one year

minus also the \$13.54 million land value). Thus, based on the NPV criterion presented in Chapter 4, including the opportunity cost of the land, the development would not currently be optimal, and the highest and best use (HBU) of the land would, at this point, be to hold the land vacant. Notice that this result is due completely to the option premium of \$1.78 million in the land value. Our analysis shows the source of the value of the land and reveals how much of this is due to the option premium. This is, essentially, an example of **the irreversibility premium** described in Chapter 7.

*16.3 A RIGOROUS MODEL OF OPTION VALUE

The previous illustration demonstrated how future uncertainty about the market for built property interacts with the irreversibility of the construction process (hence, the mutual exclusivity of building today versus building later) to give land value an option premium and make it optimal (value-maximizing) in some cases to delay development (i.e., to engage in “land speculation,” that is, holding land for investment purposes for subsequent development). However, our example analysis in the preceding section was incomplete and lacked rigor in at least one major respect that OVT will allow us to address. In particular, we needed to know the OCC of the option to wait in order to evaluate the NPV of that alternative and compare it to the alternative of building today. We *assumed* that this OCC was 20 percent (thereby reducing the expected value of waiting until next year, \$16.25 million, to a present value today of \$13.54 million). But this assumption was entirely ad hoc. We had no rigorous basis to know if it was correct, and there was no way to know what the true opportunity cost of capital of the option actually was. In this section, we will show how OVT solves this problem, and in the process, we shall discover that, in fact, our assumption of 20 percent was wrong, and hence, our evaluation of the land in the preceding section was wrong!⁶ We will present the OVT solution in two ways: (1) the **arbitrage argument** and (2) the **certainty-equivalence valuation procedure**. It turns out that these two approaches are equivalent (and we will explain why).

*16.3.1 THE ARBITRAGE PERSPECTIVE ON THE OPTION VALUE

Let's play like economists for a few minutes and make some admittedly pretty unrealistic assumptions. (It's not that we don't respect the real world; our simplifications are only meant to allow us to see the *essence* of the truth about the phenomenon we are studying.) In particular, we are going to assume the existence of something economists call “complete markets.” In fact, we will assume markets that are not only complete but “frictionless” to boot. This means that investors can buy or sell (including short sales⁷) without any trading costs, any fraction or quantity of three types of assets that are traded in three markets:

- Land
- Built properties
- Bonds

The land market trades undeveloped but developable land; the built property market trades completed, fully operational income-producing buildings (together with the land they occupy); and the bond market trades debt instruments (contractual future cash flow obligations that we will assume are riskless).

⁶ Shame on us. But we don't want to upset you by introducing too many new concepts at once.

⁷ In a “short sale,” the investor sells an asset before he buys it! (Obviously, this is the reverse of the normal, or “long,” procedure.) In order to accomplish a short sale, you have to borrow the asset you are selling. This means, of course, that you will have to buy it (or an identical asset) later in order to return it to its rightful owner (in order to “close out your short position”). In a short sale, you receive cash first when you sell the asset and pay cash later when you close out the short position by buying the asset you previously sold. It's easy to see how short sales can be arranged in the stock market (where shares of a given company are all identical). In real estate, short sales are impossible in reality, but we are going to assume for analytical purposes in this section that they are possible.

Suppose further that we can observe the opportunity cost of capital (investors' typical, expected total returns) in both the built property market and the bond market. Let's say that the OCC for riskless bonds is 3 percent per year, and the OCC for built properties like the office building that could be built on our land in the example in the previous section is 9 percent per year.

Let us return to the option valuation problem we previously considered in Section 16.2. Suppose we could purchase today a 67 percent interest in a future building a year from now, just like the building that could be built on our land for \$90 million at that time, a building that will be worth either \$113.21 million or \$78.62 million then.⁸ Such a purchase would have a price today of \$63.29 million, calculated as follows:

$$\begin{aligned}
 \$63.29 &= (0.67) PV[V_1] = (0.67) \frac{E[V_1]}{1 + OCC} \\
 &= (0.67) \frac{(0.7)\$113.21 + (0.3)\$78.62}{1 + 9\%} \\
 &= (0.67) \frac{\$102.83}{1.09} = (0.67) \$94.34
 \end{aligned}$$

We could also finance such a purchase in part by borrowing \$51.21 million today at the 3 percent riskless interest rate (given that we will certainly pay the money back next year, with interest). Thus, our net investment today would be $\$63.29 - \$51.21 = \$12.09$ million.

In effect, we would have invested in a portfolio that will mature in one year, consisting of two positions: a "long" position, which is a 67 percent share investment in the future completed building, plus a "short" (or negative) investment in a bond that will be worth for certain \$52.74 million next year (as $\$51.21 \times 1.03 = \52.74). Exhibit 16-2 depicts the situation we face. Suppose the up scenario unfolds (the market for built property improves, which you recall has a probability of 70 percent). In that case, our portfolio will be worth \$23.21 million, equal to the \$75.95 million value of our 67 percent share of the completed building (which will be worth \$113.21 million in that scenario) minus the \$52.74 million that we will have to pay back on the loan. Suppose the down scenario unfolds (the built property market falls, a scenario that has a 30 percent probability). In that case, our portfolio will be worth exactly nothing, as our 67 percent share in the building will be worth exactly the same as the balance due on our loan ($\$52.74 = 0.67 \times \78.62).

Now, compare the possible future values of our portfolio with the optimal net outcome of the development option, as described in Section 16.2. Comparing Exhibits 16-1 and 16-2, it is obvious that our portfolio will be worth exactly the same amount as the optimal net outcome of the development option, no matter which future scenario happens next year. For this reason, our portfolio is called a "**replicating portfolio**" (or a "hedge portfolio"). It replicates the option value in all possible future outcomes.

If our replicating portfolio will have exactly the same value as the option next year, no matter what, then it must have exactly the same value as the option today. We have already seen that this value is \$12.09 million (that is, we could acquire our portfolio today for \$12.09 million net cash outflow). Thus, the option to wait and develop the property next year must be worth \$12.09 million, not the \$13.54 million that we previously calculated in Section 16.2. This option value must be the value of the land, as it is worth more than the \$11.76 million immediate development value.

If the land were worth any value different than the \$12.09 million we have just determined, then an "**arbitrage**" opportunity would exist in our economist's dream world. That is, investors could earn excess profit risklessly. In effect, you could create a "money machine," as follows.

⁸ Recall our assumption (for now) that construction is instantaneous.

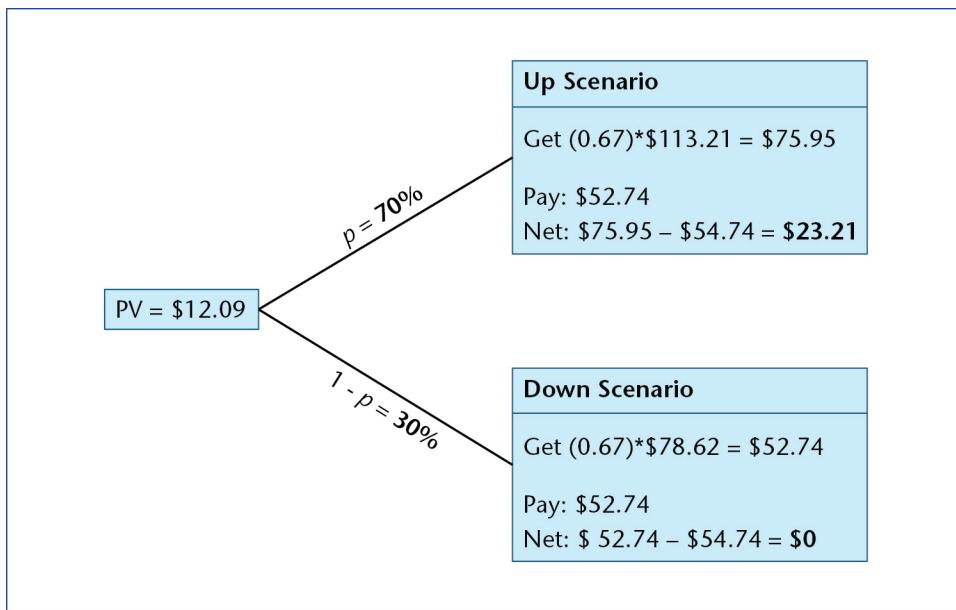


EXHIBIT 16-2 Binomial Outcome Possibilities for the One-Period Development Option Arbitrage or “Hedge Portfolio” (values in millions).

Suppose, for example, that the land price was indeed \$13.54 million. The recipe for the money machine would then be to sell the land short and purchase the replicating portfolio. We would take in \$13.54 million today from selling the land and pay \$12.09 million to acquire the replicating portfolio (as we previously described), thus netting \$1.45 million today. But what will happen next year?

Next year, we will have to buy back the land in order to close out our short position in the land, to pay back whoever had loaned us the land today a value equivalent to the value of the land next year. If the *up* scenario happens, this will require us to redeem this person with a value of \$23.21 million because that would be what the land would be worth next year in the *up* scenario. (Recall that the option expires next year, so it will definitely make sense to develop the office building at that time in the up scenario, which will produce an asset worth \$113.21 million for a construction cost of \$90 million.) Of course, this is exactly the same value as what our replicating portfolio will be worth in that scenario, so we would simply cash in our portfolio and use the proceeds to pay the original landowner, thereby closing out our short position with no net cash flow. Suppose, instead, the *down* scenario occurs. In that case, the land will be worth zero, so we can close out our short position without having to pay the original landowner anything, which is fortunate because our portfolio will also be worth nothing in that scenario. In any case, we will face zero net cash flow next year, for certain, no matter which possible scenario unfolds. Thus, the \$1.45 million we could make today is “pure gravy,” a safe profit without future liability.

Obviously, this same type of procedure could be employed to make an arbitrage profit if the current price of the land were *any* value greater than the \$12.09 million value of the replicating portfolio. Similarly, just the opposite strategy would allow us to make a riskless profit if the current land price were any value less than \$12.09 million. In that case, we would buy the land today, sell short the replicating portfolio, pocket the difference, and face zero net cash flow next year for certain. Exhibit 16-3 summarizes the arbitrage-based valuation that we have just described to evaluate

EXHIBIT 16-3

Summary of Arbitrage-Based Valuation of the Development Option

		Next Year	
	Today	Up Scenario	Down Scenario
Development Option Value	$PV[C_1] = x$ x has unknown value	$C_1 = \$113.21 - \$90 = \$23.21$	$C_1 = \$0$ (<i>don't build</i>)
$C = \text{MAX}(0, V - K)$	$x = P_0$, otherwise arbitrage		
Built Property Value	$PV[V_1] = E[V_1]/(1+0CC)$ $= [(0.7)\$113.21 + (0.3)\$78.62]/1.09$ $= \$102.83/11.09 = \94.34	$V_1 = \$113.21$	$V_1 = \$78.62$
Bond Value	$B_0 = \$51.21$	$B_1 = (1+r_f)B_0$ $B_1 = (1.03)\$51.21 = \52.74	$B_1 = (1+r_f)B_0$ $B_1 = (1.03)\$51.21 = \52.74
Replicating Portfolio	$P_0 = (N)P_{V_1} - B_0$ $P = (N)V - B$ $= (0.67)\$94.34 - \51.21 $= \$63.29 - \$51.21 = \$12.09$	$P_1 = (0.67)\$113.21 - \52.74 $= \$75.95 - \52.74 $= \$23.21$	$P_1 = (0.67)\$78.62 - \52.74 $= \$52.74 - \52.74 $= \$0$

the option to wait and develop the land next year.⁹ Assuming that in the real world, opportunities to make easy, riskless profits like this would be competed away in the marketplace, the arbitrage-based valuation of \$12.09 million must, therefore, be the option's market value.

Since the expected value of the option next year is \$16.25 million (equal to 0.7 times \$23.21 million), and the option's present value is \$12.09 million, the expected return on the investment in the option is:

$$\frac{\$16,250,000 - \$12,090,000}{\$12,090,000} = 34.42\%$$

***16.3.2 THE CERTAINTY-EQUIVALENCE DCF PERSPECTIVE ON THE OPTION VALUE¹⁰**

Alright, you say, but that economist's dream world is pretty remote from the reality of real estate markets. We don't really have short sales or frictionless asset markets for land and built properties. We could not actually implement the arbitrage described in the preceding section. While that is true, the option valuation result that we obtained is surprisingly robust. For example, we can get

⁹ A nagging question remains: how did we figure out what the replicating portfolio would consist of? The answer is easy in the binomial world (only two possible outcomes next period) we are working with here. We have two unknown quantities to define the replicating portfolio: the amount of the future completed building to purchase (label this quantity "N") and the amount to borrow (label this quantity "B"). We need the replicating portfolio to duplicate the development option outcome in each of next period's possible future outcomes. This establishes two equations that are linear in the two unknowns and which will, therefore, provide a unique solution for our two unknown quantities. Letting P represent the replicating portfolio, V represent the newly built value, C represent the development option net value, and B represent the riskless bond position, we have:

$$\begin{cases} P_1^{up} = NV_1^{up} - (1+r_f)B + C_1^{up} \\ P_1^{down} = NV_1^{down} - (1+r_f)B + C_1^{down} \end{cases} \Rightarrow \begin{cases} N = (C_1^{up} - C_1^{down})/V_1^{up} - V_1^{down} \\ B = (NV_1^{down} - C_1^{down})/(1+r_f) \end{cases}$$

In our example:

$N = (23.21 - 0)/(113.21 - 78.62) = 0.67$; $B = [0.67(78.62) - 0]/1.03 = 51.21$. (And we hope to convince you in Section 16.4 that this binomial world is not as silly as it first may seem.)

¹⁰ Stewart C. Myers "A Time-State-Preference Model of Security Valuation," *The Journal of Financial and Quantitative Analysis*, 3(1): 1-33.

exactly the same result by applying what is known as the **certainty-equivalence approach** to DCF valuation. Certainty-equivalence discounting is simply an alternative way to do DCF present value computation of risky future values or cash flows, equivalent to (but more flexible than) the traditional risk-adjusted discount rate method introduced in Chapter 4. Certainty equivalence valuation will enable the rigorous valuation of projects characterized by significant flexibility or “optionality,” such as many large-scale real estate development projects.¹¹ Such projects cannot be evaluated using traditional risk-adjusted discounting because it is impossible to know what the correct OCC to apply to the project would be. The certainty equivalence approach can often shortcut the need to know the OCC in advance, while providing a means to discover what the true OCC is.

Consider the basic element in the traditional risk-adjusted discounting approach to DCF valuation, the discounting to the present value of a single expected future value that will be obtained one period in the future, using a risk-adjusted discount rate expressed as a simple periodic return expectation;

$$PV[V_1] = \frac{E_0[V_1]}{1 + E_0[r_v]}$$

Here, we are accounting for both time and risk in the discount rate in the denominator, as $E_0[r_v] = r_f + RP_v$, where r_f is riskless and accounts for the time value of money, and RP_v is the market's required risk premium in the expected total return for the investment. But we can easily expand and algebraically manipulate this formula so that the denominator purely reflects the time value of money (the discounting is done *risklessly*), and the risk is completely and purely accounted for in the numerator:

Equation 1

$$\begin{aligned} PV[V_1] &= \frac{E_0[V_1]}{1 + E_0[r_v]} = \frac{E_0[V_1]}{1 + r_f + RP_v} (1 + r_f + RP_v) PV[V_1] = E_0[V_1] \\ (1 + r_f) PV[V_1] + (RP_v) PV[V_1] &= E_0[V_1] \\ PV[V_1] &= \frac{E_0[V_1] - (RP_v) PV[V_1]}{1 + r_f} = \frac{CEQ_0[V_1]}{1 + r_f} \end{aligned}$$

The future value in the numerator on the right-hand side, labeled $CEQ_0[V_1]$, is called the **certainty equivalent value**. Notice that the certainty equivalent value equals the (unbiased) expected value in dollars, $E_0[V_1]$, less a “risk discount” expressed in dollars, $(RP_v)PV[V_1]$. Numerically, the risk discount equals the OCC's risk premium component (a decimal or percent) times the present value of the claim in dollars. Conceptually, the certainty equivalent value is the amount such that the investment market would be indifferent between a riskless claim to receive that amount for certain and the actual claim to receive the risky amount V_1 , which could turn out to be either greater or less than $E_0[V_1]$ (given that the expectation is unbiased). Thus, the appropriate OCC to use in discounting the $CEQ_0[V_1]$ value is the risk-free rate, r_f .

Let's reconsider our simple numerical example to see how this can work in practice. We already know that our office building will be worth \$113 million (with a 70 percent probability) next year

¹¹ Another use of the certainty equivalence method is in the valuation of derivatives that are futures contracts, such as index return swaps. In a futures contract, no cash changes hands upfront, implying an equilibrium present value of zero. Obviously, no finite risk-adjusted discount rate can discount non-zero future expected cash flows to a present value of zero. However, the certainty equivalence approach can solve such valuation problems.

or \$79 million (with a 30 percent probability); revisit Exhibit 16-2 as well. Furthermore, we know the construction cost is \$90 million, the risk-free rate is 3%, and the office risk-premium is 6%. The question is, what is the present value of our option (our “right without obligation”) to build the office building next year? This is provided by:

$$PV[V_1] = \frac{E_0[V_1]}{1 + E_0[r_v]} = \frac{pV_1^{up} + (1-p)V_1^{down}}{1 + r_f + RP_v} = \frac{(0.7)\$113 + (0.3)\$79}{1 + 0.03 + 0.06} = \frac{\$103}{1.09} = \$94$$

Note that if we know the expected future value (\$103) or the future scenario for the office building (the 70 percent probability of \$113 and 30 percent probability of \$79), then it suffices to know either the OCC (in this case 9 percent) or the current value (in this case \$94) of the investment asset in order to determine the other variable (simply solve the above equation for the unknown variable), and thereby to determine the investment present value and expected return question completely. At least one of these parameters will usually be empirically observable in the investment asset marketplace. Underlying assets (that is, built properties) are relatively common and traded relatively more frequently, than vacant land. Thus, evaluating assets such as the future claim on the office building is relatively easy using risk-adjusted discounting. However, it is much more difficult to directly empirically observe either the OCC or the present value of the *option* to build the office building as we have described it.

But let us suppose (rather plausibly) that the market’s required risk premium in an investment, RP, is proportional to the “risk” in the investment as measured by the percentage spread in the possible change in the value of the investment between now and next year. In the case of the office building, this percentage spread is the difference between the +20% rise in value (to \$113 from \$94) in the “up” outcome and the -17% fall in value (to \$79 from \$94) in the “down” outcome. This outcome, a spread of 37 percent, measures the amount of “risk” in an investment in the office building. Given the office building’s required risk premium of 6 percent (equal to its 9 percent OCC minus the 3 percent risk-free interest rate), this implies that the market’s “*risk premium per unit of risk*” is 6% divided by 37%, or 0.162. This is effectively the investment market’s “*price of risk*” for this type of asset. This means that the following relationship, which elaborates on the structure we have just described, must hold:

$$\begin{aligned} \frac{9\% - 3\%}{\$113 - \$79} &= \frac{RP_v}{\frac{V_1^{up} \$ - V_1^{down} \$}{PV[V_1]}} = \frac{RP_v}{\frac{(V_1^{up} \% - V_1^{down} \%)}{(V_1^{up} \% - V_1^{down} \%)}} \\ &= \frac{9\% - 3\%}{(+20\%) - (-17\%)} = \frac{6\%}{37\%} = 0.163 \end{aligned}$$

Of course, the same relationship must also hold for the option if the market is in equilibrium. In other words, the same “*price of risk*” applies to all assets, and their required return risk premia must all have the same proportion to their risk. Labeling the present value of our option as $PV[C_1]$, and recalling from Exhibit 16-2 that the option will be worth either \$23 in the “up” outcome or zero in the “down” outcome, we must have;

Equation 2

$$\frac{RP_C}{C_1^{up} \% - C_1^{down} \%} = \frac{RP_C}{(C_1^{up} \$ - C_1^{down} \$)/PV[C_1]} = \frac{RP_C}{(\$23 - \$0)/PV[C_1]} = 0.162$$

The trouble is that, as we cannot readily observe from the market the value of either $PV[C_1]$ or of RP_c , we cannot use risk-adjusted discounting to evaluate the option. However, the above equation does allow us to assess $(RP_c)PV[C_1]$, which, using certainty equivalence valuation, is all we need in order to evaluate the option. From equation (2), we have:

Equation 3

$$\frac{RP_c}{(C_1^{up} \$ - C_1^{down} \$)/PV[C_1]} = \frac{RP_c}{\frac{\$23 - \$0}{PV[C_1]}} = 0.162$$

$$\Rightarrow RP_c = 0.162(C_1^{up} \$ - C_1^{down} \$)/PV[C_1]$$

$$\Rightarrow (RP_c)PV[C_1] = 0.162(C_1^{up} \$ - C_1^{down} \$) = 0.162(\$23 - \$0) = \$3.73$$

From equation (1), the certainty equivalence valuation formula, we have (for the option):

Equation 4

$$\begin{aligned} PV[C_1] &= \frac{CEQ_0[C_1]}{1+r_r} = \frac{E[C_1] - (RP_c)PV[C_1]}{1+r_r} = \frac{[(0.7)\$23 + (0.3)\$0] - (RP_c)PV[C_1]}{1.03} \\ &= \frac{\$16.1 - (RP_c)PV[C_1]}{1.03} \end{aligned}$$

Substituting Equation (3) into (4), we finally obtain the value of the option:

$$PV[C_1] = \frac{\$16.1 - (RP_c)PV[C_1]}{1.03} = \frac{\$16.1 - \$3.73}{1.03} = \frac{\$12.37}{1.03} = \$12$$

Having obtained the present value of the option, we can now, of course, “back out” the OCC and the risk premium for the option:

$$PV[C_1] = \frac{E_0[C_1]}{1+E_0[r_c]} = \frac{\$16.1}{1+E_0[r_c]} = \$12 \Rightarrow E_0[r_c] = \frac{\$16.1}{\$12} - 1 = 34.42\%$$

Note that this 34 percent is similar to the return on the investment in the option we found at the end of the previous Section, 16.3.1. Even though the result is similar, recall that we did not employ an arbitrage argument to derive the certainty-equivalence valuation formula. We simply assumed that, in equilibrium, assets in the investment market must trade at prices that reflect the same “price of risk” for all assets. This implies that the investment expected return risk premium per unit of investment risk must be the same for the option and the underlying asset from which it derives its value. Here, we simply apply that principle as a cross-market equilibrium condition across the markets for built property and land. The realism of this condition across those two markets is particularly strong because the land value is *derivative*, based solely on the value of the office building that can be built

on it.¹² We also assume that the relative amount of investment risk in an underlying asset and an option that depends upon that asset can be measured by the ratio of the percentage spreads in the underlying asset and the option's investment returns between the *up* and *down* possible outcomes. This is also a relatively benign assumption for derivatives and their underlying assets, as derivatives must be perfectly correlated with their underlying assets (i.e., when office buildings are worth more, the land on which they are built will be worth more).

*16.3.3 HOW THE OPTION VALUATION MODEL WORKS

Let's go back to that 34 percent opportunity cost of capital (market's expected return) for the development option that we have now computed in two different ways. (Note that this is quite different from the 20 percent ad hoc assumption we erroneously made in Section 16.2 when we misestimated the option value to be \$13.54 million.) As the risk-free interest rate is 3 percent, this implies that the risk premium in the option investment is 31.42 percent. The risk premium in investment in completed, fully operational office buildings like the one that can be built on our land is 6 percent, computed as the already-noted OCC of 9 percent for such investments minus the 3 percent risk-free rate. Thus, an investment in the development option (i.e., the land speculation) presents approximately $31.42\%/6.00\% = 5.24$ times the expected return risk premium of an investment in a stabilized office building.

How does this compare to the ratio of the risk presented by the two investments, as measured by the spreads in their possible investment return outcomes? The spread in the return outcome possibilities for investment in the future stabilized office building is 36.67 percent, from the \$94.34 million value today to either \$113.21 million or \$78.62 million next year.¹³ The spread in the return outcome possibilities for the option is:

$$\frac{\text{MAX}[0, \$113.21 - \$90.00] - \text{MAX}[0, \$78.62 - \$90.00]}{\$12,090,000} = \frac{\$23,210,000 - \$0}{\$12,090,000} = 192\%$$

Thus, the investment return risk ratio is $192\%/36.67\% = 5.24$, exactly the same as the ratio of the expected return risk premia across the two types of investment assets.

In other words, our option valuation of \$12.09 million is based on the principle that an investment in a stabilized office building and an investment in a development option that gives the right without obligation to build an office building next year both must present the same expected return risk premium *per unit of risk*. (They both present an expected return risk premium of 0.164 percent per 1 percent of investment return outcome spread: $6\%/36.67\% = 31.42\%/192\% = 0.164$.) That is, the OCC is such to provide equal risk-adjusted returns across different types of investments (developable land and built property). This is the key to equilibrium across the markets for built property and land. Otherwise, investors would tend to sell one type of investment (either land or completed buildings) and buy the other. This type of buying and selling would bid up the asset prices in the target market

¹² Recall the residual theory of land value presented in Chapter 6.

¹³ This can be computed as: $(\$113.21 - \$94.34)/\$94.34 - (\$78.62 - \$94.34)/\$94.34 = 20\% - (-16.67\%) = 36.67\%$. Alternatively, suppose an investor today pays the full \$100 million for an already existing office building, rather than only \$94.34 million for a forward claim on the building next year. Then the investor would receive next year not only the ex-dividend value of the office building, but also the net rental income it yields. As $100/94.34 = 1.06$, the implication is that the net rental payout ("dividend") is expected to be 6% of the ex-dividend value of the asset. Thus, in the "up" scenario the investor will obtain not only the asset worth \$113.21 million, but also $0.06 * \$113.21 = \6.79 million of income, for a total payoff of \$120 million. In the "down" scenario the investor will receive a total payoff of $1.06 * \$78.62 = \83.33 million. These payoffs, compared to the initial investment of \$100 million, provide the same percentage spread that we previously computed: $120\% - 83.33\% = 36.67\%$.

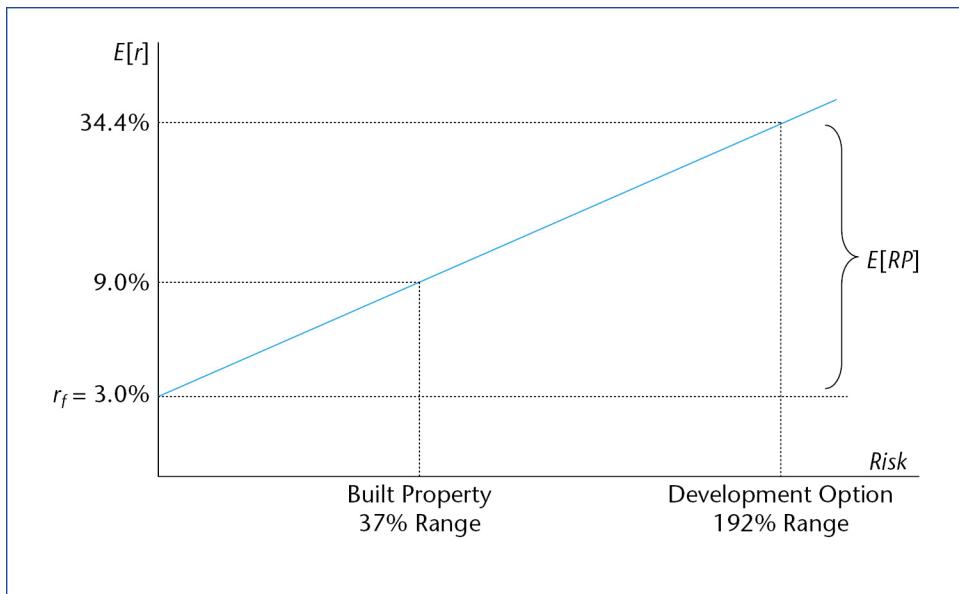


EXHIBIT 16-4 The Option Model Equates the Expected Return Risk Premium *per Unit of Risk* Across the Markets for Built Property and Developable Land.

and drive down the prices in the market being sold off, and would go on until the expected return risk premia per unit of risk were equalized across the two types of markets. The competitive drive toward this type of equilibrium across the markets is strong in a society with well-functioning land and property markets. This fundamentally makes the land valuation model described here surprisingly realistic and robust. What we are doing is pictured graphically in Exhibit 16-4.

This points out the sense in which the option valuation model may be viewed as a *normative* tool, not just as a *positive* (empirical) predictor of what the market price of the land will probably actually be. In other words, the model tells us what the “value” of the land is and what its price should be in order to provide a *fair* expected return to investors (relative to the return expectation provided by investment in built property).¹⁴ Actual prices that differ from what the model says present super-normal (or subnormal) investment return expectations either for the land or for the buildings. This is fundamentally why we get the same land valuation result with the arbitrage analysis in Section 16.3.1 as we do with the certainty-equivalence valuation in Section 16.3.2. Arbitrage opportunities, by definition, present “supernormal” profits.

Finally, note that in both the arbitrage analysis and the certainty-equivalence valuation, we were able to arrive at a solution only because we were working in a one-period binomial world in which investments can have only two possible outcomes. In the next section, we will see how this simple model can be much more realistic and useful than you might initially expect.

TEXT BOX 16.2 THE THREE MEANINGS OF OPTIONS THEORY

Students often ask how a theory, developed originally in the context of perfectly efficient, frictionless markets, can be relevant in the less “ideal” world of real estate. To gain perspective

¹⁴ Assuming one’s notion of “fairness” is that investments should provide the same ex-ante return risk premium per unit of risk in the investment.

on this question, it is useful to think of three different market contexts, or “levels,” in which option theory applications may occur in the real world. Each of these three different levels corresponds to a different implication about the meaning and use of option theory.

- **Level 1: Highly efficient markets of commodity-like assets (e.g., public stock and bond markets).** In this case, the arbitrage derivation of option theory (such as Section 16.3.1) applies directly and in reality. Suppose prices deviate from the option model prediction (assuming the model is good). In that case, you can construct a hedge that will allow you to earn supernormal profits risklessly: a true “arbitrage” in the technical sense. This situation does not perfectly exist anywhere, but it is close to existing in some aspects of the bond market. This is the part of “Wall Street” where “rocket scientists” are in great demand.
- **Level 2: Well-functioning markets but with some frictions and uniqueness of assets (e.g., most real estate markets in the United States).** In this case, the Section 16.3.2 derivation of the option model is more relevant, based simply on an argument of “equilibrium” within and across the markets for the relevant types of assets (land, built property, bonds). Equilibrium is a less stringent assumption than exercisable arbitrage. It simply assumes that markets will tend to equilibrate investors’ expected return risk premia *per unit of risk* within and across the markets that people can trade-in. In such a world, if prices deviate from what a (good) option model predicts, you may not be able to make riskless profits. Still, you can make investment decisions based on the model (buying and selling appropriately as the model suggests) so that you face supernormal expected returns (ex-ante). In well-functioning markets, prices always tend toward equilibrium, but that doesn’t mean that they will always equal the equilibrium. At this level, the option model still has what in the philosophy of science is referred to as a “positive” (or predictive) basis.
- **Level 3: No functioning market (e.g., some unique real estate situations in the United States and many situations in other countries where the market mechanism is less widely or aggressively employed than in the United States).** In this case, the option model loses its “positive” basis. However, economics is generally regarded as a “normative” science and a “positive” one. That is, it gives guidance on what “ought” to be, in some sense, not just on what “is” (in an empirical sense). Because the option model gives prices (or, at this level: “values” rather than prices, as there is no market) that equate the return risk premium per unit of risk (ex-ante), those values may be viewed as “normative,” or “fair,” in that sense. If you want the values you are working with to provide or reflect return expectations that are “fair” in the above sense, then you should use the values predicted by the option model. This could be particularly useful when entering into negotiations with other parties about how to allocate or apportion certain costs and benefits of a project or transaction that must occur in the absence of a market. It could also be useful in conducting analyses of trade-offs or “cost/benefit analysis” from a welfare or social perspective (only bearing in mind that social benefits and costs may differ from those based purely on market prices).

16.4 THE BINOMIAL MODEL OF OPTION VALUE

The one-period binomial world in which the development option examples of Sections 16.2 and 16.3 were presented is obviously a tremendous simplification of reality. In the real world, time is continuous, and asset values can assume many more than just two possible outcomes. However,

the binomial model is more than just a pedagogical simplification device. It is a building block that can be used to construct a much more general and realistic option valuation and analysis tool. In this context, it is useful to think of an individual binomial scenario (that is, a move in asset value either up or down over one period of time) as like a financial economic “molecule.” In chemistry, a molecule is the smallest, simplest particle of a chemical that still retains all of the essential characteristics of that chemical. In the case of financial economics, the essential characteristics are money (or value) across time and *risk*. The binomial element contains one unit (period) of time, and the deviation between the “up” and “down” outcomes represents the amount of risk in the particular asset being modeled.

The dynamics of the evolution of the value of built property through time, upon which the land value and the optimal development policy depend, can be modeled as a series of these molecules. Within each period, we can apply the tools presented in the preceding section for each possible future “state of the world” (up or down). By stitching the individual binomial outcomes together sequentially, we can span as long a time frame as we like, and by making each individual binomial period as short as we want, we can get the model to approach continuous time and continuous pricing realistically.

In the online supplemental materials accompanying this book, we present the binomial model of option value in detail, showing how to build realistic binomial value “trees” for the underlying asset and how to evaluate the option by working recursively backward in time within the value tree. We also provide an Excel example file online that details how the model works and can serve as a template for building your own binomial option valuation models. Here in the text, we will provide a brief idea of how the binomial model works.

First, recall our one-period world in which the underlying asset (a newly completed office building) is worth \$100 million today and next year, either \$113.21 million (70 percent probability) or \$78.62 million (30 percent probability). By reducing the period length to a month and combining 12 sequential monthly periods, we can produce the binomial value tree for the underlying asset indicated in Exhibit 16-5. As you move from left to right in the table, across the columns, you move into the future, one month at a time. Each column represents the next month in the future. Each value node in the table indicates a possible future value of the underlying asset for the development option, the value of the new office building that could be built on the land. Of course, in reality, as the future actually unfolds, only one (true) value for the building will exist at any point in time. However, the tree in Exhibit 16-5 models future possibilities as they may be envisioned from the present.

Note that the tree is constructed so that from any possible future value in any given month (each “node” in the tree), it is only possible to move to two possible future values in the next month. (You can only jump to values that are touching the right-hand corners of a given box.) Thus, the tree consists of a series of future single-period binary outcome possibilities, each one of which is conceptually just like the example one-period binomial world we considered in the previous section. We can apply the option valuation approach described in Section 16.3 to any of these one-period binomial outcome possibilities. We can quantify the option values by starting at the right-hand column, the option’s expiration, where its future values are given as the maximum of either zero or the net profit from the development project (just as in the example in Section 16.3). We can then work backward in time, from right to left across the columns, to arrive at the present value of the option.

By configuring the up and *down* outcomes and the probability of each outcome in a certain manner (as described in the online supplemental materials), we can ensure that the intermediate outcomes “recombine” as indicated in the table (the move up from one future scenario to the next column to the right results in the same value as the move *down* from the scenario immediately above it in the same column).¹⁵ This greatly simplifies the computations that must be performed.

¹⁵ For example, the \$100 million building today will be worth \$99.01 million two months from now, either because the market moved up and then down or down and then up over the intervening two months.

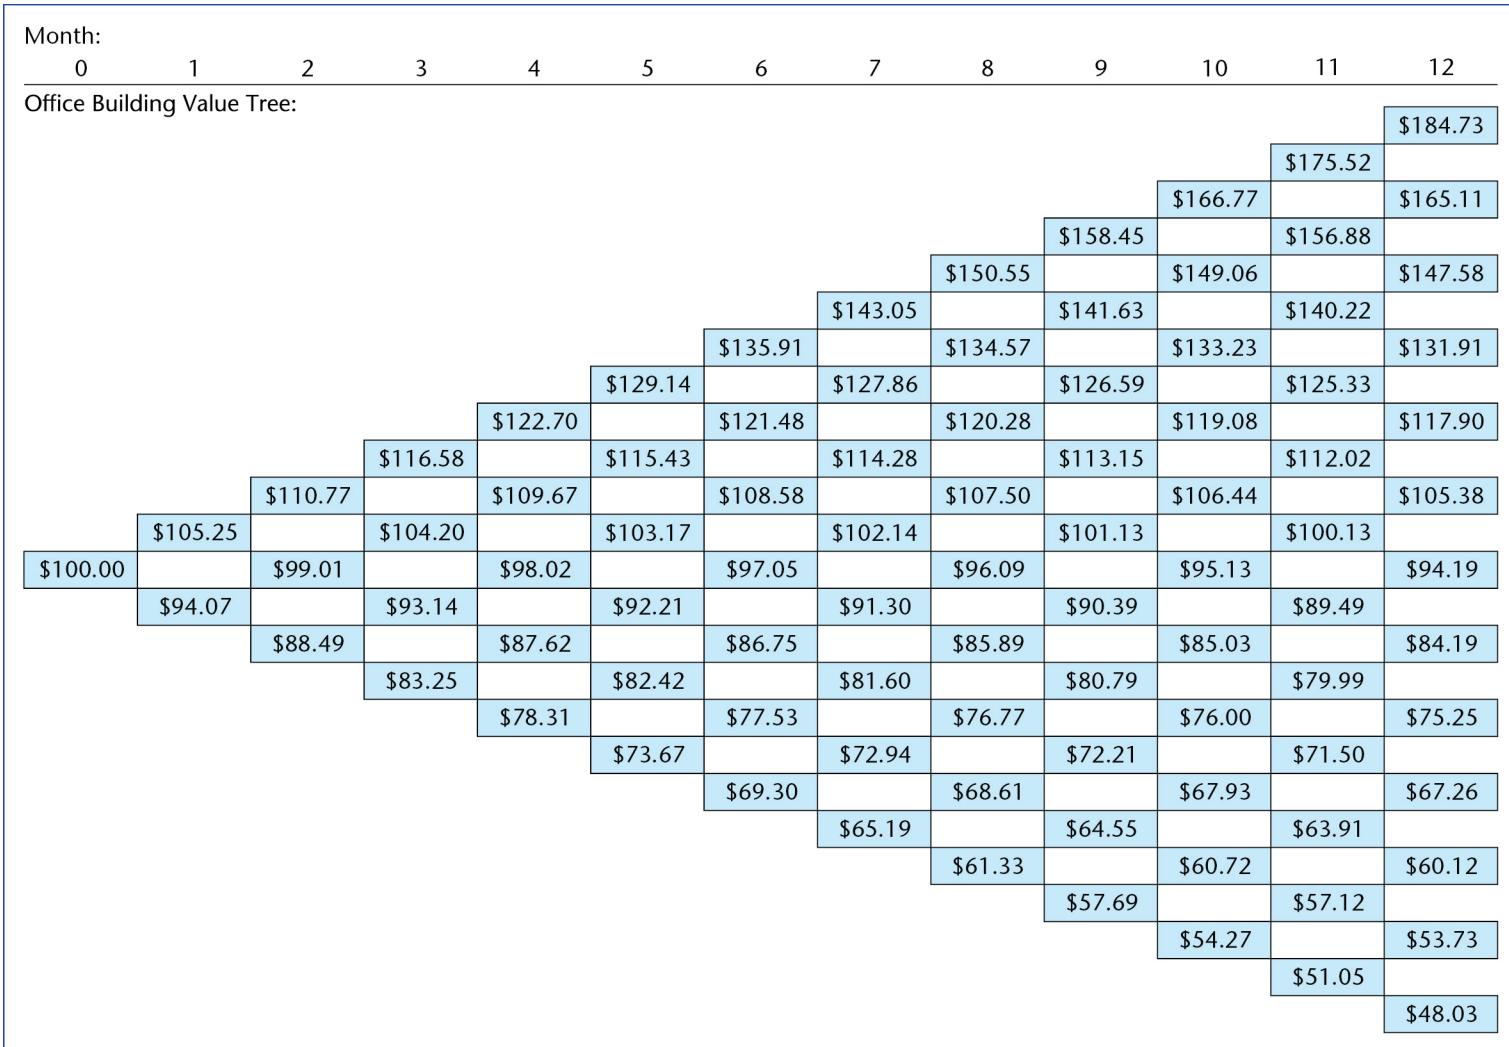


EXHIBIT 16-5 One-Year Monthly Binomial Value Tree.

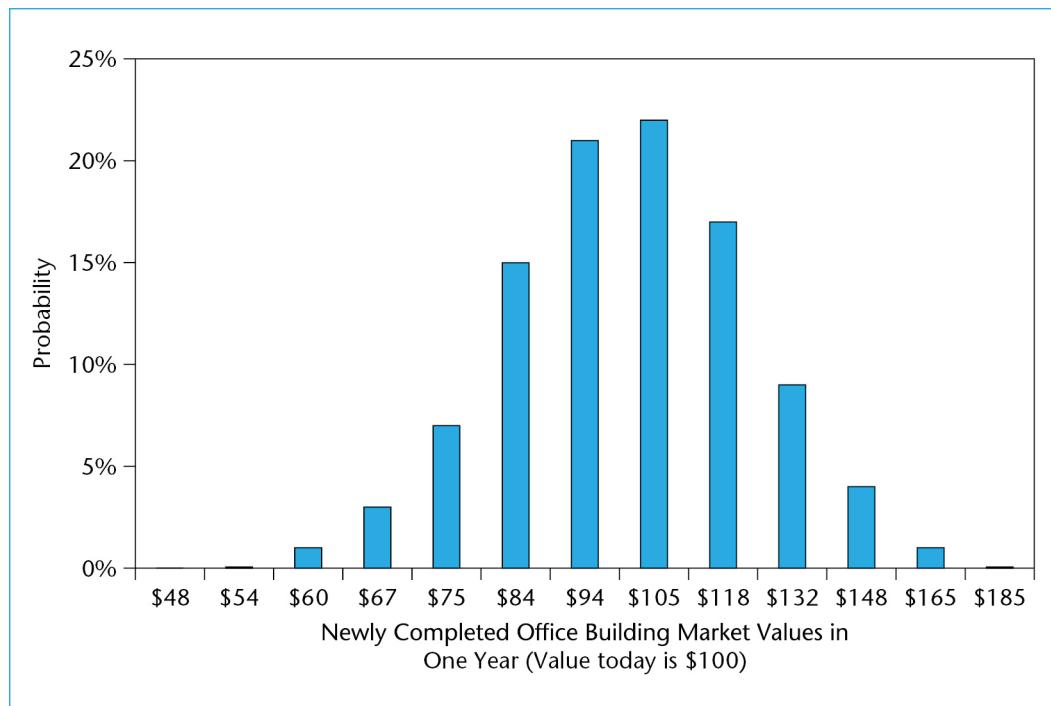


EXHIBIT 16-6 One-Year Value Probabilities for Office Building.

We can also configure the model so that the result at the end of the year is a realistically plausible value probability distribution, such as the value distribution for our office building indicated in Exhibit 16-6.¹⁶ The probability distribution indicated in Exhibit 16-6 is much more realistic than the binary possibilities of only \$113.21 million or \$78.62 million we worked with for illustrative purposes in Section 16.3. In this way, the binomial option model can be made to represent the real world well.

*16.5 A PERPETUAL MODEL IN CONTINUOUS TIME

The binomial model described in the preceding sections is the simplest and most intuitive option model relevant to land valuation and optimal development timing. It is useful not only for pedagogical purposes but also in Chapter 18, as it can be used to analyze and evaluate important real estate development project questions, such as project staging, abandonment flexibility, and development privileges that do not last forever. However, the binomial model has an important weakness for basic land valuation (as a “freehold”) and the analysis of simple as-of-right development projects that the landowner can undertake at any time. In particular, the binomial model can only be used to evaluate finite-lived options. Yet, fee simple land ownership rights are perpetual, and the right (without obligation) to build an as-of-right development project on a land parcel never expires. While perpetual development rights can be approximated in the binomial approach by modeling a

¹⁶ The probability distribution in Exhibit 16-6 has a mean of about \$103 million and a standard deviation of about ±\$20 million, consistent with an annual new-building value growth rate expectation of 3% and annual individual building volatility of about 20%. (A diversified portfolio of properties would have a lower volatility.) The model can be configured so that the binomial tree corresponds to whatever mean growth rate and annual volatility we want to input (see supplemental materials).

very long expiration time for the option, it is possible to develop exact models of perpetual options in continuous time (i.e., without the artificial discretization of time that exists in the binomial model). It will be worth our while to examine here the simplest and most elegant such model, for the insight it can give regarding the nature of land value, development timing, and the characteristics of investment risk and return in land speculation.

*16.5.1 THE CLASSICAL SAMUELSON-MCKEAN FORMULA

Our favorite continuous time option valuation formula relevant for basic land valuation is also one of the first option formulas to be developed in the economics literature. This formula, developed by the Nobel Prize-winning economist Paul Samuelson and his mathematician partner Henry McKean, was first published in 1965 as a formula for pricing “perpetual American warrants” (that is, a perpetual call option that can be exercised at any time on a dividend-paying underlying asset). This Samuelson-McKean formula turned out to be very similar to real option models subsequently developed for capital budgeting decisions by Robert McDonald and Daniel Seigel and for urban land valuation by Joseph Williams, Dennis Capozza, and others.

Although it should be treated as simplifying the complete land valuation and optimal development problem, the Samuelson-McKean formula is based very straightforwardly on the same type of economic arbitrage and equilibrium analysis described in Section 16.3.¹⁷ Thus, if you understood Section 16.3, you already understand the economic basis of the Samuelson-McKean formula. Furthermore, unlike much continuous time or perpetual option methodology, the Samuelson-McKean formula can be understood easily without specialized mathematical knowledge and can be applied easily with no more sophisticated tools than a spreadsheet. Despite its simplicity, it is a conceptually consistent land valuation and optimal development model, allowing for continuous time and an infinite time horizon.

The Samuelson-McKean formula requires as inputs three parameter values describing the underlying real estate and construction markets: the built property’s current cash yield rate, the volatility in the built property value, and a parameter we shall call the “construction cost yield.” This last is the difference between the opportunity cost of capital of construction cost cash flows and the expected growth rate in construction costs. Label the construction yield y_k , and let $y_k = r_f - g_k$, where r_f is the risk-free interest rate (which can be proxied by the yield to Treasury bills), and g_k is the expected growth rate in construction costs (which might typically be near the general inflation rate).¹⁸ Let y_v be the built property’s current cash yield rate (i.e., the net operating rent as a fraction of the property value, in effect, the cap rate of the property), which can typically be observed in the relevant asset market.¹⁹ Let σ_v be the volatility in the built property market value, measured by the standard deviation of individual property total returns across time. The relevant volatility here is that of properties

¹⁷ In continuous time, the replicating portfolio we described in Section 16.3.1 is described by partial derivatives, and the arbitrage-avoiding equilibrium condition becomes a differential equation. The solution of this equation with appropriate boundary conditions gives the Samuelson-McKean formula.

¹⁸ As described here, the Samuelson-McKean formula assumes riskless construction costs that grow deterministically at a constant rate through time. We describe how to relax the riskless construction cost assumption in Section 16A.2, located on the online supplemental materials accompanying this chapter. Note, however, that modeling construction costs as “riskless” in this context is often a good approximation of reality. For example, construction costs can often be fixed in advance contractually, and even if not, if the only uncertainty in construction costs has to do with engineering unknowns, then it is reasonable to model the OCC of construction costs as having a zero-risk premium (as the uncertainty would be uncorrelated with financial factors, and hence diversifiable). Note also that as construction costs are a negative component in the NPV of the development project, and a lower discount rate will result in a higher present value for such costs, modeling construction costs as “riskless” in this sense *magnifies* the negative impact of construction costs in evaluating the development project. This is seen in the Samuelson-McKean formula, as a lower construction yield results in a lower development option value.

¹⁹ The cap rate in this usage should be defined net of a reserve for capital expenditures so as to reflect a long-run “stabilized” yield rate for the built property.

that are already developed and in operation, not vacant land parcels, and it is the volatility of an *individual* property as distinct from that of a diversified portfolio of properties.²⁰ Typical values for this volatility measure would be between 10 percent and 25 percent per year for commercial properties in the United States (see discussion of volatility in Chapter 3).

Given values for these three parameters, we can define the **option elasticity** measure, η , by the following formula:

Equation 5a

$$\eta = \left\{ y_v - y_k + \frac{\sigma_v^2}{2} + \left[\left(y_k - y_v - \frac{\sigma_v^2}{2} \right)^2 + 2y_k \sigma_v^2 \right]^{\frac{1}{2}} \right\} / \sigma_v^2$$

For example, if $y_k = 1\%$, $y_v = 6\%$, and $\sigma_v = 20\%$, then $\eta = 3.64$. The measure η is referred to as the option elasticity because when the option is alive (not yet exercised, e.g., the land not yet developed), η gives the percentage change in the value of the option (e.g., the raw land) associated with a 1 percent change in the value of the underlying asset (e.g., built property).

Apart from the option elasticity, the other values that enter into the value of a given land parcel are the built value and construction cost of the best project that could be built on the site. Define V_0 to be the presently observable value of a newly developed property of the type that would be optimal on the site. Label as K_0 the current construction and development cost (including developer fees for service but excluding land cost and developer general overhead or sunk costs). Then, under the assumptions of the Samuelson-McKean formula, the vacant land value, here labeled C_0 , is given by:²¹

Equation 5b

$$C_0 = (V^* - K_0) \left(\frac{V_0}{V^*} \right)^\eta$$

In this formula, V^* is the **hurdle value** (sometimes called the **critical value** as well) of the developed property below which the land should be undeveloped for the time being and above which it is optimal to develop the land immediately. In other words, V^* is the current value of the completed project that signals that immediate development of the land is now optimal. This hurdle value is a simple function of the current development cost and the option elasticity defined earlier:

Equation 5c

$$V^* = \frac{K_0 \eta}{\eta - 1}$$

Thus, $\eta/(\eta-1)$ is the **hurdle benefit/cost ratio**, V^*/K_0 , that is, the ratio of built property value divided by construction cost exclusive of land cost, which triggers immediate optimal development. Note that

²⁰ Thus, this volatility includes idiosyncratic risk components that would diversify out of a large portfolio. This typically adds around 3% to the market-level volatility.

²¹ As before, we are here ignoring considerations of time to build and the time required for the construction project to be completed. In effect, we are assuming instantaneous construction. Thus, all three values in formula (2b) are valued at the same time point. Section 16A.2 in the supplemental materials describes how to account for time to build in the Samuelson-McKean formula.

this hurdle ratio is purely a function of the option elasticity, which in turn is a function of the three parameters characterizing the relevant asset markets: the built property current cash yield, the volatility of the built-property asset value, and the construction yield rate. Thus, in the Samuelson–McKean formula, the hurdle benefit/cost ratio is independent of the scale of the project (i.e., the size of the land parcel).²²

*16.5.2 GENERAL IMPLICATIONS OF THE MODEL FOR DEVELOPMENT TIMING AND LAND SPECULATION INVESTMENT RISK AND RETURN EXPECTATIONS

While the Samuelson–McKean formula may seem a bit daunting at first, it should be clear that anyone with basic spreadsheet skills could easily copy this formula into the cells of a spreadsheet and thereby harness the power of a very sophisticated economic theory to obtain some useful insights not only about the value of a given land parcel but also about the optimal timing of its development.²³ In particular, we see that the Samuelson–McKean formula not only gives the value of the land but also characterizes the construction benefit/ cost ratio (exclusive of land cost) that is necessary for immediate development to be optimal [equation (5c)]. This in itself is useful and interesting information.

Similarly, the option elasticity measure η , as given in equation (5a), is useful as a stepping-stone to derive the development hurdle ratio and land value and contains useful information in its own right. This elasticity is the percentage change in vacant land value associated with a 1 percent change in the values of built properties in the underlying real estate market. In principle, the risk premium in the expected return required by investors holding speculative land should be proportional to this elasticity, according to the following formula:

Equation 6

$$RP_c = \eta RP_v$$

where RP_c is the expected return risk premium (i.e., return over and above the risk-free interest rate) in the vacant land holding, and RP_v is the risk premium for built properties in the underlying real estate asset market.²⁴ For example, if as in our previous numerical example, the risk premium for unlevered investment in built property is 600 basis points (the 9 percent OCC minus the 3 percent risk-free interest rate), and the option elasticity in the vacant land value is $\eta = 3.64$, then the risk premium in the vacant land expected return would be 21.83% ($3.64 \times 6\%$). If the risk-free interest rate were 3 percent, then the total expected return requirement for investors would be 9 percent for built property and 24.83 percent for vacant land.

²² In online Section 16A.2.1 it is shown how V^*/K_0 is influenced by the time it takes to build the project.

²³ The major technical assumptions used in the derivation of the Samuelson–McKean formula are the following: (1) underlying asset values (in this case, built property values) are clearly observable and follow a random walk (with a constant drift rate) in continuous time; (2) instantaneous returns to underlying assets are normally distributed; and (3) the parameters in the model (volatility and yields) are known and constant. Obviously, these assumptions are violated in the real world, not only in real estate but also in other applications of real options theory. Some, but not all, of the unrealistic assumptions, can be relaxed with more sophisticated models. Nevertheless, the Samuelson–McKean formula works well in the sense that it gives results that are plausible and agree broadly with typical empirical reality. In any case, it is certainly useful for obtaining basic general insights.

²⁴ The fundamental reason for this result is that options are, formally, derivative assets based on their underlying assets. Thus, the return to the land is, in principle, perfectly correlated with the return to the built property market, only with volatility equal to η times the built property volatility. Thus, regardless of the “risk benchmark” on which the capital market bases its pricing of the risk in the built property, the land will have η times that much risk. This is, in essence, the same point we made in Section 16.3.3, where we described fundamentally how the binomial option valuation model worked, namely, that the option and its underlying asset must provide the same expected return risk premium *per unit of risk*.

The relationship between the option elasticity and the investment risk and return expectations for land as compared to built property is particularly useful because the perpetual option model is a *constant elasticity* model. That is, the elasticity does not change as a function of the underlying asset value or time (per se, holding all else constant). Furthermore, as we see in equation (5a), the option elasticity is not a function of the current “benefit /cost ratio,” V_o/K_o (This is in contrast to the case with finite-lived options.)

The relationship between built property volatility and current cash yield, land value, and the hurdle benefit/cost ratio, as implied by the Samuelson-McKean formula, are displayed graphically in Exhibits 16-7 and 16-8. In these Exhibits, the horizontal axis measures the current observable value of the underlying asset, V_o , for example, \$100 million in the case of our previous numerical example of the office building in Sections 16.2 and 16.3. The straight diagonal line represents the NPV (exclusive of land cost) of immediate development of the land (i.e., the exercise value of the development option).²⁵ The curved dashed line is the land value based on the Samuelson-McKean formula as a function of the current built property value. The point at which these two lines meet is the point at which immediate development is optimal when the built property would have value V^* . This is the first point (i.e., the lowest value of V_o) at which the development project would have a nonnegative NPV when the opportunity cost of the land (including its option premium value) is included in the NPV calculus.

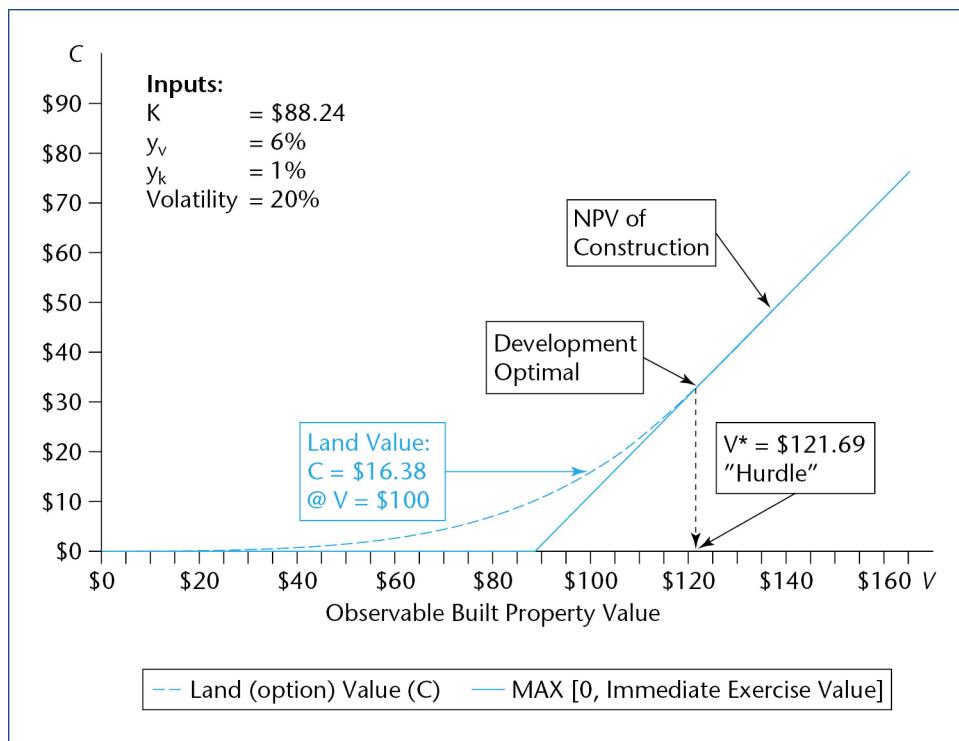


EXHIBIT 16-7 Samuelson-McKean Model Land Value as a Function of Current Built Property Value.

²⁵ Negative values of the immediate development NPV exclusive of land value are not shown, but can be inferred as the straight, diagonal line extension below the horizontal axis. Built property values less than construction cost naturally result in negative NPV of immediate development even without considering the opportunity cost of the land. On the other hand, the land itself would never be worth less than zero, as the landowner is not obligated to exercise his option.

Exhibit 16-7 is based on assumptions more or less consistent with the office building example we have been considering previously in this chapter, with current construction cost of \$88.24 million (which would be \$90 million in a year with a 2 percent inflation rate), completed property cash yield rate of 6 percent, construction yield rate of 1 percent (consistent with a 3 percent OCC and 2 percent growth rate), and completed building volatility of 20 percent. With the current built property value at \$100 million, we get an implied land value of \$16.38 million.²⁶ This is considerably greater than the \$12.09 million we obtained as the option value in Section 16.3. This is because, in Section 16.3, we assumed that the option to develop expired in only one year. Obviously, if the landowner's right (without obligation) to develop the property lasted for only one year, it would not be worth nearly as much as the perpetual option modeled by the Samuelson-McKean formula.

The Samuelson-McKean valuation of \$16.38 million includes \$4.62 million of option premium value in excess of the \$100 - \$88.24 = \$11.76 million valuation under the immediate development assumption. This implies that the immediate development of the office building would have a negative NPV of \$4.62 million, including consideration of the opportunity cost of the land. (This is found as the \$100 million benefit value of the completed built property, minus the \$88.24 million construction cost, minus the \$16.38 million value of the land.)

In this example, the hurdle value of the built property at and above which it is optimal to develop the land immediately is \$121.69 million, equal to the hurdle benefit/cost ratio of 1.3792 times the construction cost of \$88.24 million.²⁷ Only if and when we could sell the completed developed property for at least \$121.69 million should we develop the land. Notice that the Samuelson-McKean land value, given by the upward-sloping convex curved line in Exhibit 16-7, exceeds the construction NPV until the point where on the horizontal axis V_0 equals the hurdle value of \$121.69 million, in which situation the land (option) would be worth \$33.45 million, exactly equal to the construction NPV: \$121.69 - \$88.24.

For built property values in excess of the V^* threshold, if the land has still not been developed, then there is no more option premium, and the value of the land simply equals the NPV of immediate development, exclusive of land cost. However, under the assumptions of the Samuelson-McKean model, allowing the potential built property value to rise above V^* without developing the property is suboptimal, and the land should be developed immediately.

Exhibit 16-8 shows how land value and the hurdle-built property value V^* both increase with the volatility in built property values. This reveals how a greater risk in the underlying real estate market results in greater land value (other things being equal) but more tendency to delay development, as projects must pass a higher benefit/cost hurdle to make immediate development optimal.²⁸ As a result, cities with more volatile economic bases tend to have smaller, denser spatial configurations.

In fact, a similar effect (greater land value but higher development hurdle) results from a *decrease* in the built property yield rate, y_V , or from an *increase* in the construction yield rate y_K . Note that the option value with the lower volatility in Exhibit 16-8 is reduced for all values of the built property, and the hurdle value at which development is optimal is also reduced. A reduction in relevant volatility from 20 percent to 15 percent reduces the hurdle benefit/cost ratio from 1.38 to 1.22 and reduces the land value fraction of total property value at the moment of optimal development from 27.5 percent to 18 percent of the current built property value.

²⁶ We have already computed the elasticity in equation (5a) to be 3.64. Applying this elasticity value in equation (5c), we obtain a hurdle value of:

$$V^* = \$88.24[3.64/(3.64-1)] = \$121.69 \text{ million.}$$

Applying this hurdle value in equation (5b) we obtain a land value of:

$$C^* = \$121.69 - \$88.24(\$100/\$121.69)3.64 = \$16.38 \text{ million.}$$

when the current value of the built property is \$100 million (V_0).

²⁷ $1.3792 = 3.64/(3.64 - 1)$.

²⁸ Of course, increasing built property risk could also reduce the value of built properties, and this effect would reduce the value of the land.

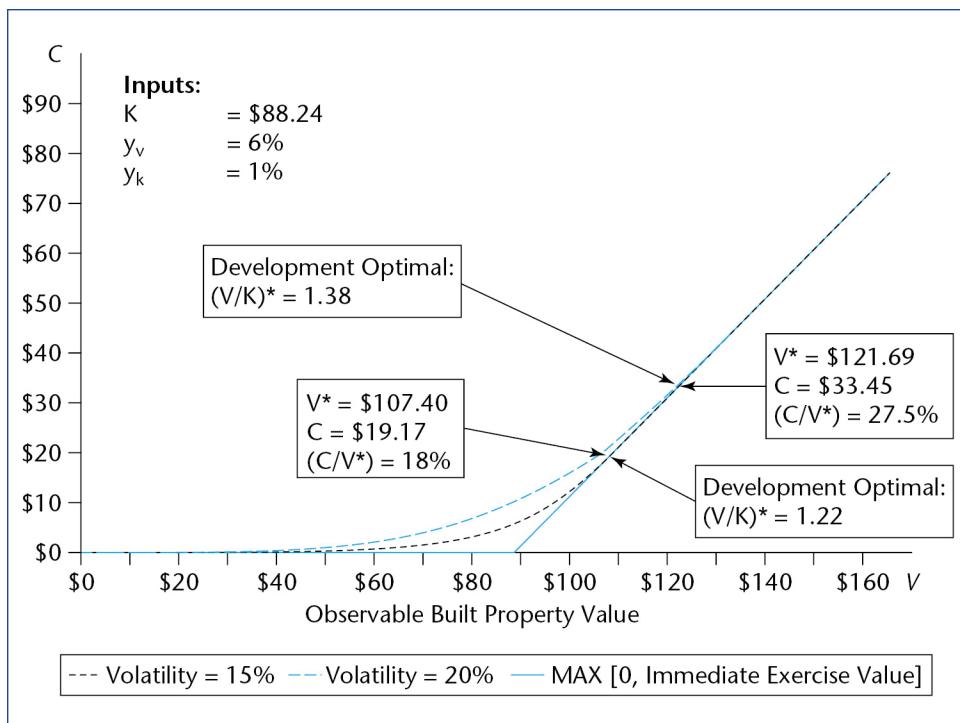


EXHIBIT 16-8 Samuelson–McKean Model Land Value as a Function of Current Built Property Value: Two Different Volatility Assumptions.

16.6 CHAPTER SUMMARY

This chapter introduced you to the concept of real options and the call option model of land value and optimal development. We also presented a simple option valuation model and the binomial model and explained the fundamentals of how and why this model works. Finally, we introduced a simple and elegant perpetual horizon, continuous time formula, the Samuelson-McKean formula, which is useful for analyzing and gaining insight regarding land value, development timing, and the opportunity cost of capital for speculative land investments. The perspective and tools presented here are useful for dealing with one of the most fundamental issues in all of real estate, whether viewed from an urban economics/space market perspective or a financial economics/asset market perspective. Land valuation and development lie at the nexus of these perspectives, at a central point in the real estate system we described way back in Chapter 6. The next chapter will step back to introduce the financial analysis of real estate development projects before we extend our treatment of the economics of land development (including real options analysis) in Chapter 18.

KEY TERMS

- Option valuation theory (OVT)
- Real options
- Option
- Exercise of options
- Underlying asset
- Exercise price
- American option

- European option
- Call option
- Call option model of land value
- Certainty-equivalence
- Value irreversibility premium
- Option premium
- Replicating portfolio
- (Or hedge portfolio) arbitrage
- Perpetual options
- Samuelson-McKean formula
- Option elasticity
- Hurdle value (critical value)
- Hurdle benefit/cost ratio
- Overbuilding
- Leasing option
- Cascades (of development)
- Development timing

STUDY QUESTIONS

Conceptual Questions

- 16.1. Define what is meant in financial economics by the concept of a call option.
- 16.2. What is a real option? How do real options differ from financial option contracts?
- 16.3. Describe the call option model of land value. What is the underlying asset in this model? What is the exercise price? What is the typical maturity of the land development option?
- 16.4. What do people mean when they say that option theory is an exception to the NPV rule? Do you agree that this is truly a valid exception to the NPV rule? Illustrate your answer with a simple real estate-related example.
- 16.5. Explain how the option model of land valuation derives from the principle of equal expected risk premiums per unit of risk across different investment assets. What are the specific asset markets that must be in equilibrium? (Hint: See Section 16.3.3.)
- 16.6. What is meant by the *normative* basis of option value theory, and how it can be meaningful even when its prescription may not be implemented or enforced by well functioning markets for the relevant assets.
- 16.7. Describe two explanations, based on real option theory, for why or how overbuilding can be due to completely *rational* (i.e., profit-maximizing) behavior on the part of developers (landowners).

Quantitative Problems

- 16.8. (Option Basics) Suppose you purchase a call option on a share of ABC Inc. common stock (i.e., go long in an option contract). This option gives you the right, but not the obligation, to buy a share of ABC's stock for \$45, which is the exercise price. Assume this is a European option, meaning it can only be exercised at the expiration date, 90 days from today. The option contract costs you \$2.50, and ABC Inc. shares are currently trading at \$40. (Note: Technically, an option contract is a contract for 100 shares, not a single share, but the economics we are concerned with are unchanged.)
 - a. Determine the current *intrinsic value* of the option defined as the payoff if the option were exercised today, where payoff is defined as the current share price less the exercise price. Explain why you would pay more than the current intrinsic value for this call option.

b. Determine what both the option payoff and your profit will be at the expiration of the contract as a function of the share price at that time (determine these quantities for share prices ranging from \$30 to \$60 at \$5 price increments). (Hint: Profit incorporates the cost of the option contract.)

c. Graph your payoff function from part (b) with share price at expiration on the horizontal axis and payoff on the vertical axis.

16.9. On a certain parcel of land, you could build a project worth \$2 million today, for a construction cost of \$1.8 million. If you wait and do not build the project until next year, the project would most likely then be worth \$2.2 million, and the construction cost would be \$1.9 million. If investors require a 25% expected return for holding land, what is the value of the land today? Should you build the project today or wait until next year?

16.10. In the preceding question, suppose there is a 50% chance that the project next year will be worth \$2.6 million and a 50% chance that it will be worth only \$1.8 million, with the construction cost still \$1.9 million in both cases. The project today would certainly be worth \$2 million and cost \$1.8 million, as before.

- Under these circumstances (and still assuming a 25% required return on land), how much is the land worth today? (Hint: Follow the approach illustrated in Exhibit 27-1.)
- Explain why the land is worth more in this problem than in the previous problem. Also explain why it is better not to build the project today, even though there is a 50% chance the project will be unprofitable next year.

16.11. (No Arbitrage/Replicating Portfolio Approach to Option Valuation) Continue with the situation in the preceding question. Assume, however, that we do not (yet) know the required return on land but we do know that the required return on built property is 8.5% and the risk-free rate of return (or opportunity cost of default-free bonds) is 4%.

- Determine the current value of the land using the no arbitrage/replicating portfolio method of section 16.3.1. Specifically, value the land as a call option by valuing a portfolio of built property and bonds that provides the exact same payoffs in the two states year 1, as the project payoffs you determined in the preceding question. (Hint: This question is similar to the example in Section 16.3.1. You should carefully work through that example to replicate key calculations. You need to determine N and B , the number of units of built property and the dollar value of bonds in the replicating portfolio, respectively, using the equations in footnote 9 of the chapter.)
- Use the certainty equivalence method of Section 16.3.2 to arrive at the same valuation.
- Determine the implied opportunity cost or required return on the land. How does your answer compare to the 25% assumed in Question 16.10?

16.12. (Samuelson-McKean) Suppose the risk-free (i.e., government bond) interest rate is 4%, the current cash yield payout rate on newly built property is 6.5%, and the annual volatility of individual property total returns is 25% for built properties that are leased up and operational. Use the Samuelson-McKean formula to answer the following questions concerning a vacant but developable land parcel.

- If built property has a 4.5% risk premium in its expected total return (8.5% total return), what is the risk premium and expected total return for the land parcel? (Hint: Use the elasticity formula (2a) and the risk premium formula (3), and note that with constant riskless construction costs the construction yield rate y_k equals the risk-free rate.)
- What is the value of the land parcel if a building currently worth \$2,200,000 new could be built on the land for a construction cost of \$1,900,000?
- What is the hurdle benefit/cost ratio above which the land should be developed immediately?

- d. What value of newly built property does this suggest is required before the land should be developed?
- e. Under these conditions, should the land be developed immediately or is it better to wait?
- f. What do you suppose is the main reason why the land value as computed here based on the Samuelson-McKean Formula is so much greater than what you computed in Question 16.11 for a parcel with similar characteristics?

16.13. (Samuelson-McKean) Suppose the risk-free (i.e., government bond) interest rate is 6%, the current cash yield payout rate on newly built property is 9%, and the annual volatility of individual property total returns is 20% for built properties that are leased up and operational. Use the Samuelson-McKean formula to answer the following questions concerning a vacant but developable land parcel.

- a. If built property has a 5% risk premium in its expected total return (11% total return), what is the risk premium and expected total return for the land parcel? (Hint: Use the elasticity formula (2a) and the risk premium formula (3), and note that with constant riskless construction costs the construction yield rate y_k equals the risk-free rate.)
- b. What is the value of the land parcel if a building currently worth \$1,000,000 new could be built on the land for a construction cost of \$800,000?
- c. What is the hurdle benefit/cost ratio above which the land should be developed immediately?
- d. What value of newly built property does this suggest is required before the land should be developed?
- e. Under these conditions, should the land be developed immediately or is it better to wait?

17 Investment Analysis of Real Estate Development Projects

Overview and Background

CHAPTER OUTLINE

- 17.1 Overview of the Development Decision-Making Process
 - 17.1.1 The Preliminary Phase of Development
 - 17.1.2 The Construction Phase of Development
 - 17.1.3 Lease-up and Tenant Finishes Phase of Development
 - 17.1.4 Finished Development: The Stabilized Asset
 - 17.1.5 Sources of Capital during Development
- 17.2 Basic Information: Enumerating Project Costs and Benefits
- 17.3 Construction Budget Mechanics
- 17.4 Evaluating Development Projects in Current Practice
 - 17.4.1 Simple Financial Feasibility Analysis (SFFA)
 - 17.4.2 Other Shortcuts to Evaluate a Development Project
 - 17.4.3 Problems with the Ad Hoc Approaches
 - 17.4.4 How Developers Think about All This
- 17.5 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The typical real estate development project decision process, at a broad-brush level, and the trend towards greener and more sustainable buildings.
- The role of financial analysis in development project decision-making and the mortgage-based simplified techniques that are widely employed in this role in current practice, including their strengths and weaknesses.
- The basics of construction loan mechanics.

We cannot overemphasize the importance of real estate development. The construction and major rehabilitation of commercial buildings is often a make-or-break activity for individual real estate entrepreneurs and firms, and it also has tremendous social and public consequences in shaping the future urban environment. Investment analysis is a crucial component of the process of real estate development decision-making. Projects that pass the financial screen get funding and get built. Projects that don't pass remain pipe dreams.

The last chapter began our in-depth study of real estate development from an investment perspective, for we saw in our study there of land value that the ultimate end of the land speculation phase is development. In this chapter and the next, we will pick up where that chapter left off by focusing on the development phase itself. In the present chapter, we will present an overview of the development process and then take an in-depth look at how development project financial feasibility is typically quantified in the real world. This chapter's practical examination of development will also briefly introduce

the mechanics of construction loans.¹ Throughout this chapter and the next, our perspective will be primarily that of the private sector for-profit real estate developer. We also focus primarily on build-to-hold development, in which the developer intends to hold the property (possibly along with partners or investors) after it is built. This model dominates in the U.S. only for commercial and rental properties. Single-family housing development is more typified by so-called “merchant builders,” who build and sell properties. However, the fundamentals of the development decision and process itself are essentially the same for all types of developers in terms of the economics of the development investment decision.

At the outset of this chapter, it is important to note that the current practice of real estate development project investment analysis relies heavily on ad hoc rules of thumb that have intuitive appeal and considerable practical value. These rules of thumb and practical procedures will be the main focus of this chapter. We will see their usefulness, but we will also see here how their ad hoc nature causes current practices to fall short in some respects, and this, in turn, will lead us to the more rigorous framework to be presented in Chapter 18.

17.1 OVERVIEW OF THE DEVELOPMENT DECISION-MAKING PROCESS

(See also 10-minute video in Chapter 17 Supplements folder)

From an economic perspective, development projects are crucial points in space and time where financial capital becomes fixed as physical capital. More broadly, they are where ideas become reality. In the preface of this book, we pointed out that real estate is a multidisciplinary field. In no real estate aspect is this more apparent and important than in the development process. In fact, development decision-making can be represented as a process that moves iteratively from one disciplinary perspective to another.

Such a development model was perhaps first articulated by James Graaskamp, a famous real estate professor and the director of the Real Estate Center at the University of Wisconsin from 1964 until his death in 1988. Graaskamp suggested that development decision-making in the private sector could typically be described by one of two situations: a site looking for a use or a use looking for a site. In the former case, the site is already under the control of the developer, and the analyst undertakes what is, in effect, a highest and best use (HBU) study. It is not uncommon for developers or land speculators to “inventory” land, that is, to buy and hold land when it is cheaper and not yet ready for development. This results in the **site-looking-for-a-use** type of decision-making, which is the type local public sector authorities are often involved in.

On the other hand, in the case of a use-looking-for-a-site, the decision maker already knows the type of development they wish to pursue. The question is, where should such a project be pursued best? What will be the cost, and what will the usage demand and competition level be at any given location? In a large development firm, the early stages of most development studies are of this **use-looking-for-a-site** type, with the developer having particular expertise in a certain type of product. A straightforward example is McDonalds (yes, they own and develop their restaurants). For them, the most important question is what location(s) are underserved. After identifying such locations, they will develop a McDonalds and search for franchisees. The use-looking-for-a-site activity may thus characterize “build-to-suit” projects and tends to be a crucial part of the business of retail firms, many of which work with real estate development firms specializing in retail development.²

The process of development analysis, design, and decision-making is highly iterative. This is depicted in Exhibit 17-1, which is based on Graaskamp’s teaching. Exhibit 17-1 is more applicable to the site-looking-for-a-use type of decision-making, which will be the major focus of Chapters 17 and 18. A given development concept will cycle through analysis from at least four different disciplinary

¹ Chapter 12 introduced permanent loans on stabilized property, and Part V will elaborate. Construction loans work rather differently in some respects, and so will be treated specifically in this chapter.

² Some developers specialize in building for a specific retailer with a fairly replicable design and finding sites that meet the retailer criteria.

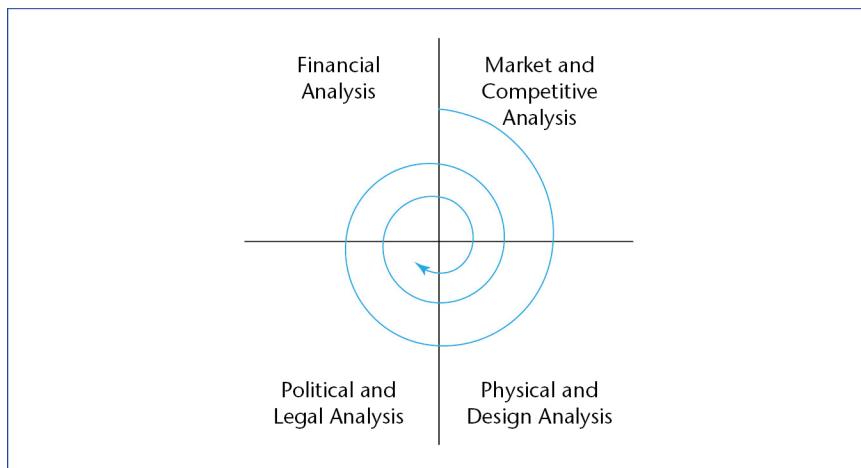


EXHIBIT 17-1 Iterative, Multidisciplinary Process of Real Estate Development Decision Making (the Graaskamp Model).

perspectives: urban economics (the real estate space market), architectural/engineering disciplines (physical analysis), legal/political analysis, and financial economics (the capital market and real estate asset market), not necessarily in that order (or indeed in any fixed order). Expertise is needed in all of these disciplines and perspectives (and sometimes others as well), and just as important, entrepreneurial creativity is needed to integrate and pull together the various perspectives to synthesize analysis from various fields into a feasible project. With each iteration, the project design and the decision become more synthesized, more detailed, and closer to fruition. This is indicated in the figure by the arrow spiraling toward the center, the point of synthesis and action.³ While all four disciplines and perspectives portrayed in Exhibit 17-1 are important, our focus in this chapter and the next is on the top left quadrant of the picture, the financial analysis of the project from an investment perspective.

From this perspective, Exhibits 17.2 and 17.3 provide an introductory overview of the development project investment process, including the major sources of capital. In these exhibits, the horizontal axis represents time, progressing from left to right. The solid line referenced to the vertical axis on the left-hand side represents the cumulative total amount of financial capital that has been expended on the project up to each point in time on the horizontal axis. In Exhibit 17-2, the dashed line referenced to the right-hand axis represents the degree of risk investors face at each point in time, viewed as the probability of the project's financial or economic "failure." Notice that the risk declines over time as the project progresses and overcomes challenges, and as it comes closer to fruition, more is learned about key factors that determine the project's outcome, such as the nature of the rental market for the project. Notice also that risk declines as more capital is invested (or, alternatively, more capital can be attracted to the project as its risk declines). The exhibit's timeline on the horizontal axis can be divided into four phases. Next, we will detail each phase.

17.1.1 THE PRELIMINARY PHASE OF DEVELOPMENT

The first phase, at the left of the horizontal axis, may be viewed as a *preliminary phase*. This is the most creative and entrepreneurial phase in the project's evolution. In most projects, this phase will be led by an entrepreneurial development firm (at least in the United States).⁴ In the case of many large, complex

³ Most projects will involve only a cursory review of all four quadrants and be dropped without further research. Each iteration requires more time and money and is pursued when the expected benefits exceed the cost of the next iteration.

⁴ In recent decades a growing number of these development firms are organized as publicly traded REITs or REOCs. Nevertheless, private development firms, both large and small, still dominate the entrepreneurial development field in the United States, and in the early 2000s the growth of the private equity sector in the capital market facilitated continued private ownership. However, it is not uncommon for private, local firms with specialized expertise and local knowledge and

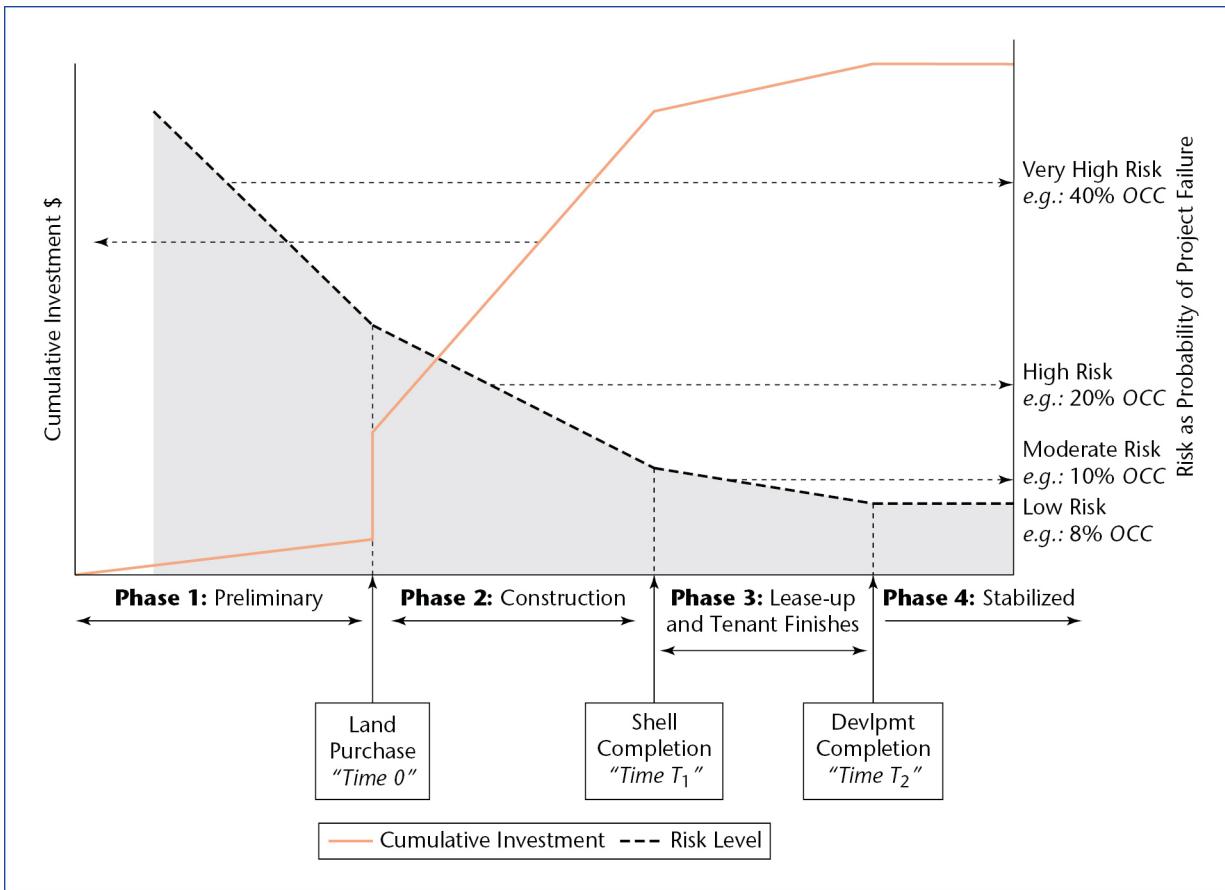


EXHIBIT 17-2 Development Project Phases: Typical Cumulative Capital Investment Profile and Investment Risk Regimes.

projects, this entrepreneurial role will be played by a governmental entity or a partnership among local government, one or more private developers, and perhaps a non-real estate firm or institution. (Such public/private partnerships are especially likely in large urban areas.⁵) Even in fairly small-scale and routine development projects, this first phase often involves the optioning and assembly of separate land parcels, obtaining necessary permits (sometimes including variances or other special provisions or infrastructure requirements), and the design of the development project. These tasks will typically proceed in parallel and iteratively with the assembly and permitting process. They will include (implicitly, if not explicitly) a “highest and best use analysis” of the site. This includes a particularly important aspect of the development from both an economic and urbanistic perspective: the optimal *density* of the development. In practical quantitative terms, density is defined in development by the **Floor Area Ratio (FAR)**, which is the building floor area divided by the land area.⁶ (See textbox.) It is in this preliminary phase of the development process that the iterative design process described in Exhibit 17-1 is most relevant.

TEXT BOX 17.1 DETERMINING THE OPTIMAL FAR FOR YOUR DEVELOPMENT

Recall from Chapter 6 that Highest and Best Use (HBU) dictates the property type for each plot when developing. Indeed, the landowner will normally only sell to the highest bidder, which depends on the optimal property use for that location and the cost to develop the relevant property type. The HBU also reflects the density that is best for the site (recall Section 7.3.1). The FAR is therefore dictated by the HBU as much as vice versa. Thus, in principle (viewed purely from an economic perspective), the developer knows how dense to develop, and simultaneously what is the maximum price they can bid for the land, by determining the FAR that maximizes land value.

To understand what maximizes land value, note that, as a general rule, as you increase the FAR, the price you can get for the property tends to *decrease*, even as the costs of development tend to *increase*, in both cases *per building square foot*. There are various reasons for this, related to construction costs, building efficiency (net rentable floor area divided by gross building floor area), and the value per square foot that tenants are willing to pay in rent.⁷ However, this cost of density is offset by the obvious economic advantage of density, which is the greater productivity of the land. The more rentable square feet per land square foot (greater FAR), the greater the value of the land site, other things equal. As a result, there is a trade-off and an optimal FAR that maximizes the land value. There is an elegant model of such optimal density based on economic land value maximization, which has been elaborated by Wheaton & DiPasquale (1996), which we will present very briefly in this text box.

Let P be the expected building value and K be the construction cost per building gross square foot, as a function of different values of FAR. The residual price (or value) per square

connections to partner with larger national firms or funds (public or private) who bring access to the capital markets and national or international connections and expertise.

⁵ For example, a university, hospital, or research institution may partner with real estate developers in conjunction with city planners. In recent years, major developments in the infotech and biotech fields have been catalyzed and located adjacent to preexisting university and medical facilities. In other cases, the establishment or expansion of a major industrial or service firm, or the development or redevelopment of a sports facility or transportation terminal, has been the generator and nucleus of complex adjacent development projects.

⁶ For example, a FAR of 2 means that the building has twice as much square footage as the land square feet of the site, such as, if the building footprint fully fills the land plot and the building has two stories.

⁷ To a point, FAR can be increased without increasing building height by increasing the proportion of the land site covered by the building footprint. But while this may not increase construction cost per built square foot, it will often reduce the rent per square foot that tenants are willing to pay.

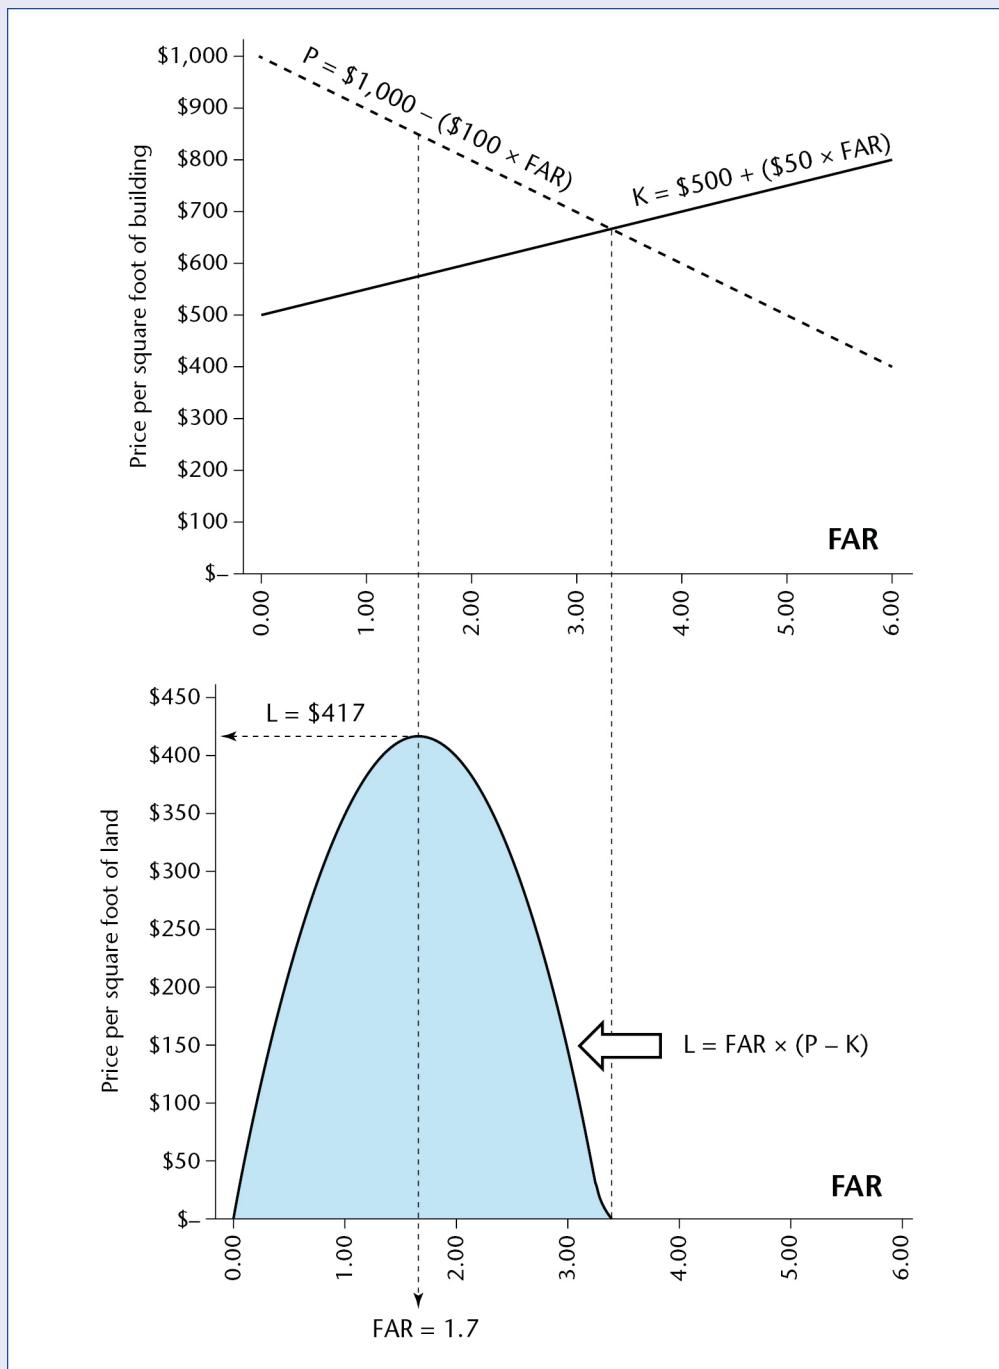


EXHIBIT 17-3 Calculating the Optimal FAR.

foot of the building is thus $P-K$. To find the price (value) per square foot of *land*, we multiply $P-K$ times the FAR: $L = FAR \times (P-K)$. Because both P and K are functions of the FAR, the result is a land value function that is parabolically shaped over the FAR, as shown in the exhibit. In the top panel, we see that the property value per built square foot (P) decreases and the costs (K) increase as the FAR increases (horizontal axis), using simple linear demand and supply functions for illustration.⁸ The top boundary of the shaded area in the bottom panel represents the resulting land value per square foot of land ($L = FAR \times (P-K)$) over the same FAR. This land value parabola has a maximum of \$417 per land SF at a FAR of 1.70. This is, therefore, the optimal density, reflecting the HBU of the site, the maximum economic bid price for the land.

The preliminary phase may take anywhere from a few months to more than a decade! Value creation, if any, occurs because of skills in envisioning and designing the project (based on an understanding of the market potential), getting through the land assembly and entitlement process, adept market analysis and targeting, tenant relationships and marketing, and a host of other skills which go far beyond just economics (which are the focus of this book). Furthermore, the preliminary phase does not always succeed. The development project may ultimately not be approved or may not prove economically, financially, or administratively feasible. In that case, much, if not all, of the entrepreneurial developer's time and expenses incurred in the preliminary phase may well have been wasted. This is why this is the riskiest phase in the development process, as indicated by the height of the dashed line in Exhibit 17-4. Referencing the right-hand vertical axis, we see that such high risk carries with it a commensurately high opportunity cost of capital (OCC) in the form of a high expected return for any financial capital invested.

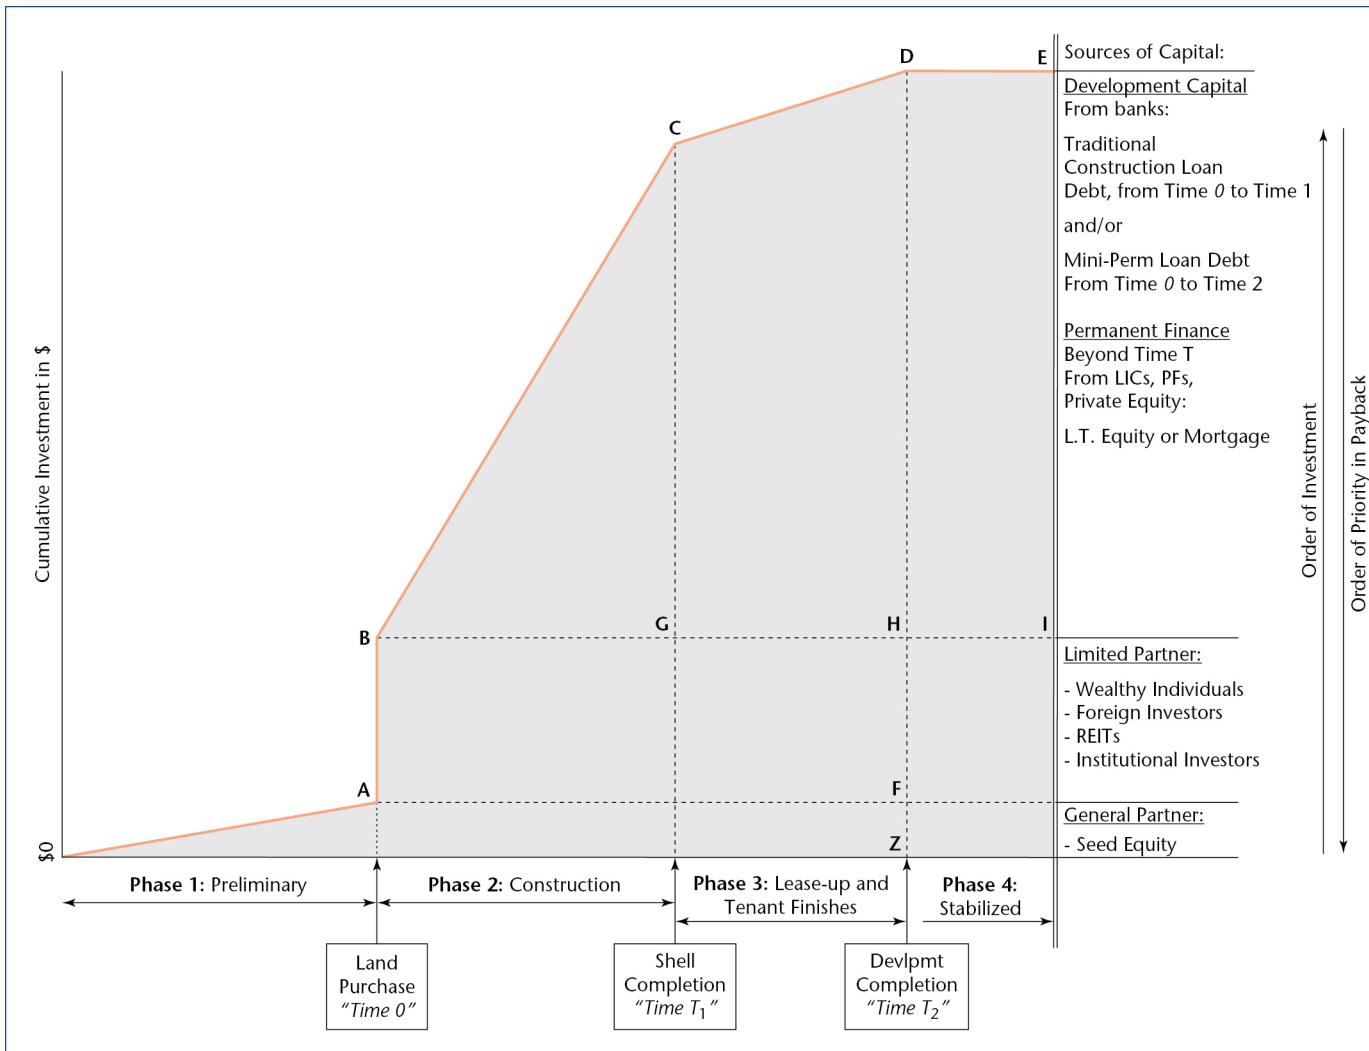
It should be noted that the amount of financial capital invested during the preliminary phase is normally relatively small *per plot*. Most of the resources spent on project development during this phase are better characterized as normal operational expenses and time during which the overhead of a development firm accrues rather than as investments of financial capital.⁹ Many of these expenses may be recouped in fees charged by the developer during or at the end of the entrepreneurial phase. However, the development firm will research many such plots, most of which will not end up being developed. Thus, the total costs for the development firm on a *firm level* are much higher, and any successful development must cover the costs of all the failed proposals.¹⁰

Profit (including return on any invested financial capital) will be obtained via the increase in land value that may be created during this phase. To the extent that financial capital has been invested in land acquisition during the preliminary phase, this may be viewed as "land speculation" investment rather than as real estate "development project" investment, and such land speculation investment may be recovered at the end of the preliminary phase if the land is sold to a separate development venture entity. As discussed in Chapter 16, land speculation investment is indeed highly risky and may be modeled using real options theory. Here in Chapter 17 (and in Chapter 18), our main focus picks up at the end of any such land speculation phase.

⁸ For illustration, in the exhibit, we assume a starting sales price of \$1,000 per square foot, which decreases \$100 for every unit increase in FAR. ($P = \$1,000 - \$100 \times FAR$.) For the construction costs, we assume a \$500 base cost per square foot which increases \$50 per unit increase in FAR. ($K = \$500 + \$50 \times FAR$.) In practice of course, these price and cost functions would be specified based on the particulars: property type, location, design etc. You can find an Excel file online, which allows you to play with the inputs.

⁹ Recall that the land may not be actually acquired during this phase but rather may only be "optioned" (as described above), or it may already be owned by the developer (perhaps purchased previously as a land speculation investment or otherwise acquired for various reasons, or inherited).

¹⁰ Anecdotal evidence tells us that only one is developed of every 10 plots seriously considered by large-scale REITs. And that is out of about 100 plots that were on the initial "longlist," of which 90 were quickly discarded.



TEXT BOX 17.2 HOW DO DEVELOPERS EXTRACT COMPENSATION?

Most of the compensation to developers is aligned with and after passive investors (limited partners or shareholders) receive priority minimum returns. For example, the limited partners may receive an 8% preferred return, after which the split is 20% to the general partner (developer) and 80% to the limited partners. The developer may receive a stepped-up basis in the land value and have this counted as part of their capital contribution. For example, the land cost \$10 million two years ago, but now that entitlement has been approved, it is revalued at \$25 million and entered as the developer's capital contribution. There may be a fee for the acquisition of a site or building to be retrofitted, typically 1 to 2%. There may be asset management or property management fees, which will most often align with market-based fees. There may be vendor fees to firms owned by the developer, such as a golf course management company or a cleaning company, all of which provide extra compensation to the developer. These fees may all be market-based, but since they are under the control of the developer, they should all be transparent, disclosed, and hopefully market-based.

At the end of the preliminary phase depicted on the horizontal axis of Exhibit 17-4 (if not before), the necessary land will be purchased or otherwise acquired to give the developer the right to proceed with construction.¹¹ Regardless of when and how the land has actually been acquired, from the perspective of the evaluation of the development project as an investment, the economic "opportunity cost" of the value of the land is incurred at the time of the start of construction. We label this point "Time 0" in Exhibit 17-4. At this point, the land is irreversibly committed to the construction project, and it ceases to be a "land speculation" with real options characteristics. It is important to recognize, however, that the economic opportunity cost of the land at Time 0 may differ from the historical monetary cost of land acquisition. The opportunity cost is the economic value of the land, what it could be sold for at Time 0 after it has been assembled, permitted, and a project design developed. The entrepreneurial developer may have created much of this land value during the first phase of the overall development process depicted in the graph. Thus, the actual price paid for the land, represented by the solid line in Exhibit 17-2, may be quite different from its economic opportunity cost (hopefully, from the developer's perspective, much less).

17.1.2 THE CONSTRUCTION PHASE OF DEVELOPMENT

The development investment per se actually begins with the incurring of the land opportunity cost at the beginning of the *construction phase*. It is during the construction phase that the bulk of the financial expenditures are usually required, for the purpose of paying for the construction of the building(s) on the site.¹² This phase is still quite risky, though less so than the preliminary phase. Development investments are risky for two reasons. First, if the project is "speculative" (referred to as building "on spec"), then it is not known for certain what occupancy will be achieved within what period of time and at what rental rate. This is fundamental rental market risk, as described in Part II of this text, and in the context of the development project, it can be said that spec development is exposed to "lease-up risk."

The second reason development investments are risky is that they inherently contain "**operational leverage**." Even if there is no financial leverage (that is, if the project is financed entirely with

¹¹ This would normally include the exercise of any land purchase options, as described above. (To the extent that some or all of the preexisting landowners are not being paid off at this time, they should be viewed as development partners.)

¹² In some high-value locations, land costs may approach or even exceed construction costs in total magnitude.

equity), and even if there is no lease-up risk (that is, if the project is entirely preleased), development projects still have operational leverage in the sense of having high fixed or committed costs relative to potentially variable revenues. This will be discussed in depth in Chapter 18. Still, the essence of the matter is that the development investor incurs the land opportunity cost at Time 0 and commits to paying the construction costs at the end of the construction phase, no matter what the value of the project turns out to be at the time of completion, which we label “Time T” in the exhibits.

For example, suppose the opportunity cost of the land at Time 0 is \$20 million, and the construction costs during the development phase are \$70 million, with the expectation as of Time 0 that the value of the project upon completion at Time T will be \$100 million. Thus, \$20 million is invested at Time 0 with the expectation of \$30 million being returned at Time T (the expected difference between project value minus construction cost: \$100–\$70 million). However, the construction cost of \$70 million must be paid no matter what. Suppose the project turns out to be worth \$90 million instead of \$100 million (only a 10 percent decline in value). In that case, the net at Time T is cut from \$30 million to \$20 million, and the return on the \$20 million Time 0 investment is reduced to zero. Because of this type of operational leverage, the construction phase of development investments requires a rather high expected return, reflecting an OCC often in the neighborhood of 20 percent per annum (going-in IRR for the development phase only).

Interestingly, recent research has shown that *without* such operational leverage, development is actually *less* risky compared to stabilized assets at the right of Exhibit 17-2.¹³ The reason is that developments provide some flexibility. In one extreme example, you could postpone the construction phase entirely to “ride out” a market lull. In other words, you have some discretion in the timing, even after you purchase the land.¹⁴ You can also tweak certain design choices as market circumstances change. Note that the earlier you are in the construction cycle, the more flexible you are. If you are nearly done developing the property, not much can be changed (cheaply) anymore. Once your property is fully developed, you are fully subject to market forces and will not be able to benefit from such flexibilities.

17.1.3 LEASE-UP AND TENANT FINISHES PHASE OF DEVELOPMENT

The next phase on the horizontal axis in Exhibit 17-2 is described as “*Lease-up & Tenant Finishes*.” This is the phase after the major construction of the “shell” has been completed, during which its users lease and occupy the space. For many types of commercial development where long-term leases are involved, this involves some finishing construction work that customizes the space(s) in the building for the tenant(s) who will be occupying the space. While the tenants may pay for some of these expenditures, it is not uncommon for the developer to pay much or all of these expenditures. Nevertheless, this phase usually involves less capital and less risk than the main construction phase of the development project.¹⁵ As a result, investment in the lease-up phase involves a lower OCC, for example, perhaps around a 10 percent required rate of return (as suggested in the exhibit).

The construction and lease-up phases together represent the “development project” per se, the subject of the investment analysis in this chapter. These development phases may take anywhere from a few months to several years. In larger projects (with multiple buildings), the development

¹³ See, Geltner D., A. Kumar, and A. Van de Minne, 2020, “Riskiness of Real Estate Development: A Perspective from Urban Economics and Option Value Theory,” *Real Estate Economics* 48(2).

¹⁴ We should note that in some countries like China, developers are not permitted to sit on undeveloped land indefinitely. Approved developments must be initiated within a rather short timetable or totally abandoned.

¹⁵ With nonspeculative (preleased) development, the lease-up phase may be effectively nonexistent, with the tenant finishes phase absorbed in the main construction phase. At the other extreme, in speculative development of building types that involve major tenant-specific construction, such as biotech space and some types of retail space, the Lease-up & Tenant Finishes phase may involve substantial capital requirement and risk. (For this reason, developers usually try to avoid speculative development of such space.)

may often be divided into separate stages, in which there remains flexibility about when (or even whether) to commence the later stages of the development. In such cases, the development project retains real options characteristics, as defined in Chapter 16.

17.1.4 FINISHED DEVELOPMENT: THE STABILIZED ASSET

At the end of the development project, the result is an asset in “*stabilized operation*.” The project is completely or nearly leased up and operating at its long-run steady-state level of profitability (at least relative to the condition of its rental market). This point is labeled Time 2 in Exhibit 17-2 (or just Time T for short) to collapse lease-up and construction into a single “development” phase. At Time T, the development phase investment is complete, and what exists going forward is an investment in a stabilized operating asset. Usually, there is a recapitalization at this point because different types of investors typically are in the market for stabilized (low-risk) investments. If the construction was financed by a “construction loan” (as it usually is), this loan will be retired (or “taken out”), either by an equity infusion or (more commonly in the case of taxable investors) by taking out a “permanent” commercial mortgage on the property.

The equity investment in the project may or may not be recapitalized at this time by means of a sale of the asset. However, whether or not the original equity investors sell the stabilized asset, the risk and return nature of their investment has changed. It now has the nature of the real estate acquisition investments that were the subject of Parts I–III of this book, and the investment from Time T going forward should be analyzed and evaluated as such. It is important in the modern investment world not to “mix apples and oranges.” Investment characteristics and performance should be compared among like-kind investments. Therefore, a rigorous framework for the analysis of development investments should take the Time T moment of development completion as the ending horizon of the analysis.

17.1.5 SOURCES OF CAPITAL DURING DEVELOPMENT

Exhibit 17-4 maps the typical financial capital sources and types onto the previous development project cumulative investment chart. As depicted to the right of the right-hand vertical axis in the chart, three major types of financial capital are typically provided throughout the overall development project (including the Preliminary Phase). (Follow the “Order of Investment” arrow in Exhibit 17-4 leading up to the capital sources list on the right.) The first to come in is the entrepreneurial developer’s (i.e., the General Partner, from Chapter 15) seed equity investment, which is invested primarily during the Preliminary Phase and is normally very small in absolute dollar magnitude compared to the other, later sources of capital. The next investment capital to come in is usually external equity, financing the land purchase. This external equity requirement is typically much greater than the previous seed equity requirement. It may come from one or more of a variety of sources (as well as the entrepreneurial developer itself), including opportunity funds, financial institutions (such as pension funds, life insurance companies, or endowment funds), foreign investors, wealthy individuals, private equity vehicles,¹⁶ or in some cases, REITs as joint venture partners with the original developer.¹⁷ It is not uncommon for this “external equity” component to include more than one type or class of investor and investment, often in a more or less complex partnership structure, as described in Chapter 15 (what we referred to as limited partners, see Section 15.3). It is also not uncommon for some of the capital in this component to come in the form of subordinated debt, such as “mezzanine

¹⁶ This would include, among others, private REITs and tenant-in-common (TIC) vehicles designed to funnel small individual investor capital into real estate investments.

¹⁷ Of course, in some cases, the REIT will be the “general partner” themselves. In fact, REITs do a lot of development themselves.

debt.” Finally, the last capital to come in is typically the bulk of the total capital requirement, in the form of a “construction loan,” typically from a commercial bank or syndication of such banks (for larger projects). The order of investor payback seniority is typically the reverse of the order of investment but corresponds to the magnitude of capital invested and the inverse order of the amount of investment risk (riskier investments tend to get paid back last).

The original entrepreneurial developer’s seed equity investment is represented in the OAFZ region of Exhibit 17-4. This capital comes in first (during the preliminary phase) and typically takes the highest risk and highest residual return position in the overall capital structure of the project. This is appropriate, as the entrepreneurial developer’s capital was invested (if at all) during the highest-risk (“preliminary”) phase, as we described previously in our discussion of Exhibit 17-2.¹⁸ Apart from fee income that may partially or entirely reimburse the developer for operating and overhead expenses, the entrepreneurial developer’s return on capital invested will usually be subordinated to that of later and larger capital providers. However, the entrepreneurial developer will retain a disproportionate share of any excess return available after subsequent senior capital providers have been paid their preferred returns. Given the entrepreneurial developer’s normally small investment (in absolute dollar terms), the result can be very large percentage returns when the project is successful. (For example, it would not be uncommon for an entrepreneurial investor to have an expected return in the neighborhood of the 40 percent OCC suggested in Exhibit 17-2 for capital invested during the Preliminary Phase.) This structure provides the entrepreneurial investor with the maximum incentive to see the project through to successful completion. This is an appropriate and important incentive structure since the project’s success normally depends most heavily on the developer’s expertise, ability, and effort.

The ABIF region of Exhibit 17-5 represents the external equity investors’ investment (i.e., the “limited partner”). As noted, this capital normally pays largely for the land purchase, which may or may not include some profit to the entrepreneurial developer. This purchase happens at Time 0 and kicks off the development project per se. This equity investment is normally necessary to secure debt financing through a construction loan. Construction loans normally require first liens on the property, and they require some preexisting equity investment to provide a degree of default (credit loss) protection to the lender.

The external equity investment (possibly including some subordinated debt vehicles such as “mezz debt”) is a high-risk investment that seeks high returns if the project is successful. However, such investors will typically accept a less-than-proportional payment of residual returns (after principal payback, relative to the entrepreneurial developer) in return for seniority over the entrepreneurial developer in earning “preferred returns,” which are typically intermediate in value.

For example, a typical arrangement might be as follows. Upon project completion, the construction lender is paid first, providing perhaps a 6 percent return to the construction loan (assuming no default). The remaining available value in the project then accrues to the limited and then general partner *pari passu* (proportional to their investments) until their investment principal is returned, or, depending on the arrangement, perhaps up until both partners earn a prespecified preferred return. In some cases, the external equity investors (LP) might be paid on a preferred basis until they have reached a prespecified hurdle. Preferred returns might be on the order of 10 percent or so, depending on how risky the project appeared to be at the time of the partnership agreement and how much external capital was needed by the entrepreneurial developer (and also depending on the seniority structure in the arrangement and on how the remaining residual is to be split). Finally, any remaining value in the project after preferred returns are paid is split in a nonproportional way, with the entrepreneurial developer (GP) receiving more than her pro rata share (earning a “promote”)

¹⁸ Alternatively, the entrepreneurial developer may have elected to take all or a portion of her entrepreneurial (Preliminary Phase) investment profits in the form of land sale profit at *Time 0*, by selling the land into the development partnership that would normally be created at that time to include the new, external equity investor partners.

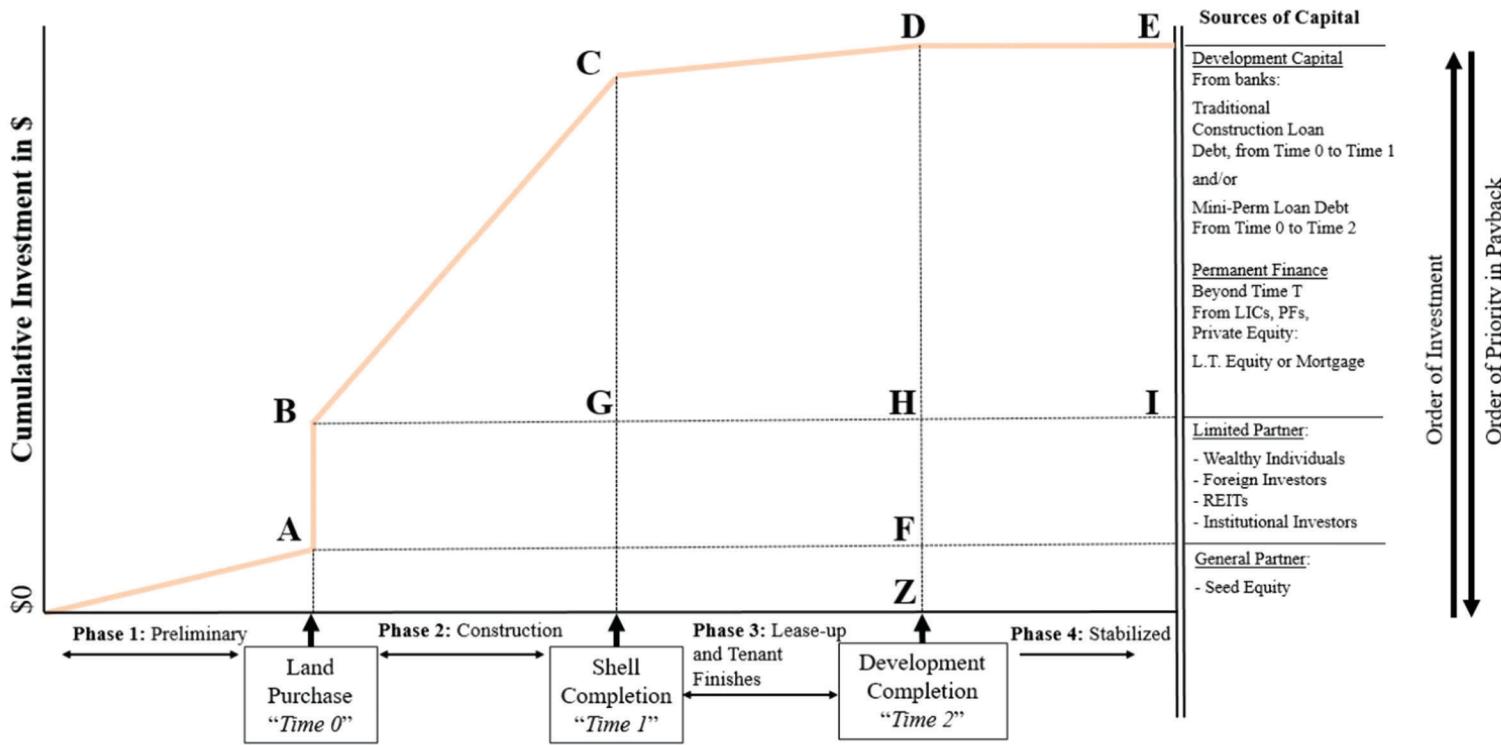


EXHIBIT 17-5 Development Project: Typical Sources of Investment Capital.

but still providing the external equity investors with some considerable additional upside potential. (Exactly how much more varies from case to case.) Including this expected residual component, the external equity might be looking for expected returns in the neighborhood of 20 percent (as suggested by the OCC in Exhibit 17-4 based on the time when the external equity capital came in), but possibly with an asymmetric rightward skew in the return probability if the preferred return has seniority.¹⁹ As noted in Chapter 15, there are infinite specific variations on this general framework of capital structure for the equity portion of the development project's financing.

As noted, the bulk of the actual development capital will typically come in the form of a first-lien construction loan. A classical construction loan will cover the BCG triangle in Exhibit 17-5, covering only (but often all of the) construction costs to the physical completion of the project. If the project has a lease-up phase and requires working capital²⁰ during that period subsequent to construction completion but prior to stabilization, a "bridge loan" may be necessary to cover the GCDH region. Alternatively, a so-called "mini-perm" construction loan may be arranged at the outset, which covers construction costs plus working capital through to the stabilization of the property (region BCDH in the exhibit).

TEXT BOX 17.3 THE TREND TOWARDS GREENER DEVELOPMENT

The term **sustainable development** refers to real estate development projects with physical characteristics that promote environmental and social improvement. The term "green" is also used to describe sustainable development, though sustainability is not restricted just to environmental considerations; it includes factors such as historical preservation and social and community enhancement. In the twenty-first century, green development is especially focused on energy efficiency and the closely related concern about the "carbon footprint" of the construction and operation of the project (considering the entire life cycle, including potential effects of ultimate demolition).

Development has become increasingly green, typically defined as either reaching Energy Star label certification, LEED (Leadership in Energy and Environmental Design) certification, and other labels discussed in Chapter 28. Of course, a fundamental question is whether the costs to develop or retrofit buildings with more sustainable features that save energy or reduce carbon impact are worthwhile. In fact, the new reality is that for investors and developers who are thinking about the long run (if they are at all considering the obsolescence factor in the structures they are building), there is little choice today but to build to a higher standard. This is true for all property usage types. The market standards for all property types are constantly being redefined, with local building codes requiring more energy and water-efficient buildings and many tenants insisting upon such higher standards, such as the federal General Services Administration (GSA).

Research has shown that tenants want three features more than any others: natural light, temperature control, and good ventilation. This applies to all property types, but especially office and retail. Whether new development or retrofits, we find the following are common upgrades:

¹⁹ If the external equity is given seniority in the claim on the preferred return, then there will be ex-ante a return probability "spike" at the preferred return rate (for example, 10%), with very small probabilities of earning substantially *more* than the ex-ante expectation (for example, 20%). The result is a rightward-skewed return probability distribution. (Though there is also some possibility of earning *less than* the preferred return, including, in the worst case, a total loss of the equity investment.) However, be careful about taking the 20% target noted here too literally. Development investment returns can be defined in different ways. (Also see our discussion in Chapter 18.)

²⁰ Working capital refers to operating cash flow deficits that occur when the building is incurring operating expenses (probably including marketing and leasing expenses) not yet fully offset by rental income revenue.

- Reduced plug loads by monitoring occupancy and automatically turning things off when not in use.
- Lighting with LEDs and fluorescents, and occupancy sensors that turn off or adjust lighting to the natural light available
- Windows that open; ventilation that is more occupant-controllable
- Cooling with chilled beams, under-floor vents, natural air, and screening out the sun's heat by design
- Heating passively with the sun by design, and more efficient fuels such as natural gas fuel cells or solar concentrators
- Wind turbines, PV solar cells, reduced water flow features, natural landscaping, and using captured rainwater and gray water (water from the building sinks) for irrigation.
- Better insulation and vented wall designs with better glass glazing that reflects heat
- More use of local and recycled materials as well as renewable wood stocks
- No use of toxic chemicals in carpets, glues, or wood products.

Commissioning a building is the process by which developers improve the operation of heating, cooling, lighting, and ventilation systems so that the building is comfortable, safe, and efficient. This is becoming standard practice as well.

Upon stabilization (full or nearly full steady-state occupancy), the construction loan will be paid off using so-called “permanent financing,” aka the “Takeout.” This can be obtained through either long-term equity and/or a long-term commercial mortgage. The mortgage may come either from a conduit CMBS lender or from a whole loan portfolio lender (as described in Part V of this book). Such long-term investment in stabilized assets, whether equity or debt, normally comes from relatively conservative institutional sources, such as pension funds, life insurance companies, as well as endowment funds, foreign investors, and wealthy private individuals. On an unlevered (property-level) basis, the stabilized asset investor might be looking for an 8 percent return (as suggested by the OCC in Exhibit 17-4) or even less for institutional-quality property. (See Chapter 11 as well for some empirical and recent OCC data.)

With this overview of the overall real estate development process in mind, we will now focus the remainder of this chapter on the investment analysis and evaluation of the central step in this process, the construction phase described here, including the commitment of the land to the development project. The commitment of the land and the construction includes the vast bulk of the total capital expenditure on the development project. In focusing on this phase of the process, however, we will keep in mind both the preliminary phase and the subsequent lease-up phase, and we will relate some of our discussion to these phases as well.

17.2 BASIC INFORMATION: ENUMERATING PROJECT COSTS AND BENEFITS

In the iterative process depicted in Exhibit 17-1, one of the information products that needs to be refined progressively is a realistic budget for the project. In fact, two types of budgets are important in development projects, each focusing on a different time period or phase of the project. The first is called the **construction and absorption budget**. It covers the period from the beginning of construction until the building is fully leased and operational. It relates to the investment cost side of the NPV equation. It includes the cost of land acquisition (or opportunity cost of land development), site preparation, hard and soft construction and design costs, and lease-up costs, including the need

for working capital until the building breaks even.²¹ The second type of budget is known as the operating budget. It covers the period beyond the end of the absorption or lease-up of the new structure. The projected operating budget underlies the benefit side of the NPV equation, namely, the value of the newly completed property. As we know from Parts I–III of this book, this value is based on the projected net cash flows the completed building is expected to generate.

In analyzing development projects, the operating budget and the resulting estimate of completed project value are often focused on a single so-called **stabilized** year, with a valuation based on the direct capitalization of that year's NOI rather than on a multiyear DCF analysis. Such a budget is referred to as a stabilized annual operating budget, as it applies to a point at which the building's cash flow has become more or less stable, reflecting its long-run ongoing operating characteristics. As such, the operating budget is typically characterized by items you became familiar with in Chapter 5:

- Potential gross income (PGI)
- Vacancy allowance
- Effective gross income (EGI)
- Operating expenses (and capital reserve)
- Net operating income (NOI)

The development of the operating budget is derived from an analysis of the relevant space market, the upper-right quadrant in Exhibit 17-1, based fundamentally on the principles of urban economics and market analysis methods.²²

The other side of the NPV equation, the cost side, dealt with in the construction and absorption budget, is a very different animal. It arises from the lower-right-hand quadrant in Exhibit 17-1, based fundamentally on the engineering and architectural disciplines. The major component of the construction and absorption budget typically concerns the construction phase. Costs associated with this phase are traditionally divided into two major categories: **hard costs** and **soft costs**. The former includes the direct cost of the physical components of the construction project: building materials and labor. Soft costs typically include just about everything else in the construction phase, such as design, legal, and financing costs. (The latter includes the interest on the construction loan.) Land acquisition or opportunity cost is sometimes classified as a soft cost, although it is more commonly treated as a hard cost or as a separate item altogether. Typical hard and soft cost items are shown here:

In addition to the construction phase, many development projects will also require the developer to budget for an **absorption or lease-up phase**, as indicated by the last items in the previous list. An absorption budget will be necessary when the project is being built at least partially “**on spec**” (short for “on speculation”); in other words, at least some of the space in the project is not preleased at the time the development decision is being made. The developer must line up in advance the financial resources necessary to carry the project through to the break-even point. In addition to the working capital necessary to operate a less-than-full building, the absorption phase budget often requires particular expenditures for marketing the new building. The absorption phase may include major expenditures on leasing commissions and tenant improvements (TI) or **build-outs**.²³

From a property-level cash flow perspective, the absorption phase may be considered to be over when the new building begins to break even on a current cash flow basis. However, from a financial perspective, the absorption phase continues until the building is stabilized (as this term was defined

²¹ There are a variety of construction cost data sources that may help in the preliminary analysis including but not limited to RSMeans Data: Construction Cost Estimating Software and Design Cost Data - Design Cost Data and Dodge Construction Network | Construction Projects and Bidding

²² A detailed exposition of this quadrant is beyond the scope of this book, although a basic introduction to space market analysis was provided in Part II, Chapter 6.

²³ Build-outs (or tenant finish) are tenant improvements on unfinished space, for example, to finish and customize space for a specific long-term tenant.

EXHIBIT 17-6

Typical Subdivision of Construction Costs

Hard Cost	Soft Cost
Land cost*	Loan fees
Site preparation cost	Construction interest
Shell costs of existing structure in rehab projects	Legal fees
Permits	Soil testing
Contractor fees	Environmental studies
Materials	Land planner fees
Labor	Architectural fees
Equipment rental	Engineering fees
Tenant finish	Marketing costs, including advertisements
Developer fees	Leasing or sales commissions

earlier, that is, until the building is at or near its expected long-term occupancy level). Only when the stabilization point is reached does the investment risk in the asset fall to that of a fully operational, ongoing property, and only then is the permanent lender typically willing to provide a traditional long-term commercial mortgage.

17.3 CONSTRUCTION BUDGET MECHANICS

A key feature of development projects is the fact that construction takes place over a period of time, typically several months to over a year, depending on the scale and complexity of the project. Both in theory and in practice, the effect of the temporal spread of construction costs is often dealt with through the device of a **construction loan**, which enables the developer to avoid most cash outflows until the project is completed. At this time, the construction loan must be paid back, including interest that has accrued on the funds drawn down during the construction project. In this section, we will review the basic mechanics of construction loans.

When a construction loan is initially signed, it normally provides the developer with a specified maximum amount of cash in the form of **future advances** based on the projected construction budget and schedule. Funds are **drawn down** out of this commitment as they are needed to pay for the construction put in place. The lender verifies the physical construction before disbursing each **draw**. In the classical construction loan, the developer/borrower typically pays no cash back to the bank until the project is completed.²⁴ At that point, the developer may take out a **permanent loan** secured by the built property and use the proceeds to repay the construction lender. Of course, there are other ways to pay off the construction loan, as the developer may invest his own equity, obtain long-term equity partners, or sell the property.

During the construction period, the developer will draw down the construction loan based on a schedule agreed to at the outset. For smaller projects, invoices and receipts from subcontractors, along with affidavits that all prior work has been completed and paid for, will usually suffice for the lender, along with some site inspections, to be sure that the work and materials specified in the budget have actually been used. For large projects, cash might be drawn down from the loan commitment based on a completion schedule of construction phases, such as site preparation, foundation work,

²⁴ This practice varies in different areas as for projects of different sizes, and also on the reputation of the developer and the relationship the developer has with the lending institution. In recent years, and especially for large loans, it is not uncommon for interest-only payments to be required on the current balance on an ongoing basis, even during the construction process.

framing, and so on, with each component resulting in the draw for a certain percentage of the total loan. Even though the construction budget typically includes contingency amounts for unexpected costs, lenders often hold back a reserve in the range of 10 percent of the total construction budget. This reserve will only be paid out when the project is completed, as evidenced by a **certification of occupancy** from the local building inspection authority.

At the time the construction loan is being negotiated, the lender will often require an engineering review of the proposed construction budget. At that time, the construction lender will also typically consider the projected lease-up phase and review the projected stabilized operating budget of the completed building from the perspective of permanent loan underwriting criteria, as described in Part V of this book. The construction lender must feel confident that the completed project will support a permanent loan sufficient to pay off the construction loan.²⁵ As mentioned earlier, a permanent loan that is used to pay off a construction loan is often referred to as a **take-out loan**. Indeed, the classical method of construction lending is for the institution making the construction loan to require that the developer obtain a commitment in advance by a permanent lender sufficient to cover the projected construction loan balance due. The permanent lender considers the projected loan/value (LTV) ratio and debt service coverage ratio (DCR) forecasted for the completed project before committing to the permanent loan. The permanent loan commitment is often contingent on the building achieving a certain occupancy or rent level by a certain time.²⁶

In order to develop the construction budget and estimate the amount of cash the developer will have to come up with at the time of project completion (or the amount he must borrow in the take-out loan), a key calculation needs to be performed of the **accrued interest** that will be due when the construction loan is paid off. For example, for a construction loan with an 8 percent interest rate (per annum, monthly compounding), the accrual of interest in the loan balance is illustrated in the following table for the first three months of a construction project, in which a total of \$2.75 million of direct construction cost has been expended and financed by the loan. Note that the loan balance exceeds this direct cost amount (\$2,780,100 is greater than \$2,750,000). This is the effect of the time value of money (including a risk premium for the construction lender), as unpaid accrued interest is added to the balance and compounded forward across time.²⁷

In the example in the table, the interest at the end of the first month is computed as $\$3,333.33 = \$500,000(0.08/12)$. This accrued interest then becomes part of the loan balance on which further interest will accrue. The interest computation in the second month is thus $\$8,355.55 = (\$503,333.33 + \$750,000)(0.08/12)$. The balance at the end of the second month is computed as $\$1,261,688.88 = \$503,333.33 + \$750,000 + \$8,355.55$. This type of accrual and compounding of interest continues until the construction project is complete. The total balance due on the loan then includes both the direct construction costs and the financing cost, or the cost of construction capital, namely, the interest on the funds borrowed for construction. This interest cost is listed in the typical construction budget as a separate item. See online supplements for a spreadsheet example of a construction loan.

²⁵ Note that, as the permanent loan will typically not be made for more than a 75% LTV ratio, this implies that the construction costs must generally be less than the total value of the completed project. The difference represents the developer's equity in the project, which often equates more or less to the cost of the land.

²⁶ Sometimes the lease-up phase is covered in the construction loan (a "mini-perm" loan), and sometimes it is covered in a separate "bridge loan."

²⁷ In the example in the table, it is assumed that the draws are made at the *beginning* of each month, and the loan balance refers to the balance at the *end* of the month. This is the most common approach for budgeting, although end-of-month draws or beginning-of-month balances are sometimes used. For budgeting purposes, it does not matter which assumption is made as long as the requisite interest is accrued as soon as funds are disbursed from the bank. (This is an application of the "four rules" of loan payment and balance computation first described in Chapter 12.)

EXHIBIT 17-7

Example of the Mechanics of a Construction Loan

Month	New Draw	Current Interest	New Loan Balance
1	\$500,000	\$3,333	\$503,333
2	\$750,000	\$8,356	\$1,261,689
3	\$1,500,000	\$18,411	\$2,780,100
4	and so on		

17.4 EVALUATING DEVELOPMENT PROJECTS IN CURRENT PRACTICE

So, how does the investment decision-making process in real estate development projects go? Chapter 4 taught you how to apply the NPV decision rule when investing in stabilized property. You are right if you think the NPV decision rule should also apply to real estate developments! Indeed, though you may not want to admit it (and we would hate to think we've done it to you), you may actually be "thinking like an economist." After all, the NPV investment analysis and evaluation model is nothing if not based squarely on fundamental economic principles, including opportunity cost, market equilibrium, and wealth maximization; principles that not only are fundamental in economic theory but reflect eminent common sense in a capitalist society.

Unfortunately, in the real world, investors barely use the NPV decision rule *explicitly*. Therefore, before we go into great detail on how it *should* be applied to development investments, in Chapter 18, we will lay out a few current tools real estate developers do use in current practice (Sections 17.4.1 and 17.4.2). Obviously, there should be value for students to understand how real-world practice works. However, this will also allow us to introduce a more complete and correct framework by explaining the shortcomings of the common methodologies from the perspective of rigorous and proper economic theory (Section 17.4.3).

17.4.1 SIMPLE FINANCIAL FEASIBILITY ANALYSIS (SFFA)

We turn first to a widely used methodology that often takes the place of the investment evaluation of development projects, but in fact doesn't even claim to *evaluate* projects, but rather merely to assess their financial *feasibility*. Even in common parlance, "feasibility" is not the same thing as "desirability." But if you're already sure that you would like to do a project, if you could, then financial feasibility is an obvious analysis to do. We call this approach **Simple Financial Feasibility Analysis (SFFA)**.

Let's see how it works by considering some concrete examples. For this purpose, we can call on Bob, a typical real estate developer on Main Street. Bob owns several strategically located, commercially zoned vacant land parcels in Midwest City and specializes in building small to medium-sized office and retail buildings. Some of his projects can be characterized as a site-looking-for-a-use, while others are better described as a use-looking-for-a-site. The iterative approach depicted in Exhibit 17-1 describes his overall project design and decision-making process well. Bob employs a simple, easy-to-understand methodology to deal with the financial analysis aspect of this process (the upper-left-hand quadrant in Exhibit 17-1).

The basic idea in SFFA is simple. It is assumed that the developer will "borrow to the hilt," that is, he will take out the maximum permanent loan the completed project will support. (After all, this is what Bob usually does.) With this assumption, Bob does not need to know (or assume) anything about the capital market except what he can easily observe in the commercial mortgage market. Through his contacts among local mortgage brokers and bankers, Bob can easily keep track of current mortgage interest rates and underwriting criteria, such as the maximum allowable LTV ratio and DCR requirements. Combined with the borrow-to-the-hilt assumption, this information is sufficient to ascertain either what sort of rental market is required given a land and construction cost

or what sort of land and/or construction cost can be afforded given a rental market for a completed structure. In other words, Bob can apply SFFA to address the financial feasibility of both the site-looking-for-a-use and the use-looking-for-a-site types of projects.

Take the site-looking-for-a-use first. Bob's procedure for analyzing financial feasibility often starts with a presumed land and construction cost and ends with the required minimum rent per square foot of built structure that will be necessary to make the project financially feasible.²⁸ The specific steps in this procedure are shown in Exhibit 17-8.

A good example of this procedure in action was a rehab opportunity Bob considered recently. A major owner/occupant had recently vacated a certain class B office building in a fantastic location. The building was 30,000 gross SF with 27,200 SF of leasable area. Bob knew that rents for class C+ or B- space in that location were running around \$10/SF per year, with most operating expenses passed through to the tenants.²⁹

Bob was pretty sure he could buy the building for \$660,000 and fix it up for about \$400,000 in hard costs and \$180,000 in soft costs, including funds for a new marketing program, financing during rehab, and operating costs during the lease-up period, which he estimated would take only a few months. A local lender indicated they would provide an 80 percent LTV mortgage with a 120 percent DCR if Bob personally signed the note for 20 percent of the loan balance. The interest rate was 11.5 percent on a fixed-rate 20-year amortization, monthly payment mortgage. Bob's market research

EXHIBIT 17-8

SFFA Front-Door Procedure

Site acquisition costs
 + Construction cost
 = *Total expected development cost*

Loan-to-Value
 = *Permanent mortgage*

× Annualized mortgage constant
 = *Cash required for debt service*

× Lender-required Debt-Coverage-Ratio
 = *Required Net-Operating-Income*

+ Expected landlords' operating expenses
 = *Required Effective Gross Income*

+ Expected occupancy rate
 = *Required gross revenue*

+ Leasable square foot
 = *Rent required per square foot*

Order of calculations



²⁸ When the developer already controls the site, he typically has a fairly firm idea of how much the land is worth. The land value, combined with legal/political and architectural/engineering considerations, then provides a general implication for the usage type and the scale or magnitude of the project that probably makes sense. (In theory, the Optimal Density model described earlier in this chapter is also relevant.) This suggests the rough magnitude of construction cost that makes sense. (And in principle, the hurdle benefit/cost ratio described in Section 16.5 of Chapter 16 is relevant in this regard, suggesting the magnitude of construction that would make immediate development optimal.) In the iterative analysis, design, and decision-making process characteristic of development projects, this cost information then begs to be checked against what the demand in the space market will support.

²⁹ Note the development analysis procedure's iterative nature or circularity. Even though this is a site-looking-for-a-use, Bob already has a rough idea of the use, at least in terms of the typical market rents.

indicated that less than 5 percent vacancy was typical for well-located, newly rehabbed projects. Given the superior location of this building, Bob figured a 95 percent stabilized occupancy rate was, if anything, conservative. Based on prior occupant figures, Bob estimated stabilized operating expenses at \$113,000 annually. Even though some of this might be passed through to the tenants, Bob decided to assume initially that the landlord would bear all of these expenses in order to be conservative, and because much of this expense figure was based on an old (but efficient) centralized heating system that made energy costs difficult to prorate to individual tenants. The question Bob wanted to get a quick take on was, "What are the required rents for this project, and, therefore, does it look feasible?"

Applying the procedure in Exhibit 17-9, here is how Bob answered his question:

Financial feasibility requires a gross rent of a little over \$10/SF. As Bob's assumptions in applying the previous procedure were (he believed) a bit on the conservative side, and market rents did indeed appear to be around \$10/SF, the project looked feasible to Bob.³⁰

The SFFA procedure applied in the previous manner is sometimes referred to as **front-door feasibility analysis**. A **back-door feasibility analysis** would go in the reverse order, starting from expected market rents and ending with the amount that can be afforded for site acquisition. Back-door analyses are more typical of use-looking-for-a-site type projects. The back-door version of the SFFA is presented in Exhibit 17-10.

An example application of the back-door version of SFFA was another recent project Bob considered. In this case, the project was a new development, an office building with 35,000 SF

EXHIBIT 17-9

Bob's Example Calculations in the SFFA Front-Door Procedure

Site and shell costs:	\$660,000
+ Rehab costs:	\$580,000
= Total costs (assumed equal to value):	\$1,240,000
× Lender-required LTV	80 percent
= Permanent mortgage amount:	\$992,000
× Annualized mortgage constant:*	0.127972
= Cash required for debt service:	\$126,948
× Lender-required DCR:	1.2
= Required NOI:	\$152,338
+ Estimated operating expenses (landlord):	\$113,000
= Required EGI:	\$265,338
+ Projected occupancy (1-vac):	0.95
= Required PGI:	\$279,303
+ Rentable area:	27,200 SF
= Required rent/SF:	\$ 10.27/SF

*From Chapter 2, we have the monthly mortgage constant: δ

$$(0.115/12)/[1 - 1/(1 + 0.115/12)^{240}] = 0.010664.$$

This figure times 12 is the annual constant: $0.010664 \times 12 = 0.127972$.

³⁰ Note that to apply this procedure Bob had to assume that the value of the completed project as it would be judged or appraised by the lender would be at least equal to his total projected construction cost of \$1.24 million. His familiarity with the local-property-asset market suggested that typical cap rates for similar buildings once fully rented seemed to be around 10.5%. Applying this to the projected NOI of \$152,338 implied a projected property value of \$1.451 million, well over the \$1.24 million construction cost, so Bob felt he was safe in this regard.

EXHIBIT 17-10

SFFA Back-Door Procedure

Total leasable square foot*
 × Expected average rent per square foot
 = *Project Potential Gross Income*
 - Vacancy allowance
 = *Expected Effective Gross Income*
 - Projected landlords' operating expenses
 = *Expected Net Operating Income*
 ÷ Debt Service Coverage Ratio
 ÷ Annualized mortgage constant
 ÷ Maximum loan-to-value ratio
 = *Maximum supportable total project costs*

Q: Can it be built for this?

-Expected construction costs (other than site)
 = *Maximum supportable site acquisition cost*

Q: Can we acquire the site for this or less?

* equals the building efficiency ratio times the gross area

Order of calculations

EXHIBIT 17-11

Bob's Example Calculations in the SFFA Back-Door Procedure

Potential gross revenue (29,750 x \$12):	\$357,000
- Vacancy at 8%:	\$28,560
= Effective gross income:	\$328,440
- Operating expenses not passed on:	\$63,000
= Net operating income:	\$265,440
÷ 1.20 (required debt coverage):	\$221,200
÷ 12 (monthly debt service):	\$18,433
⇒ Supportable mortgage amount:	\$2,048,735
÷ 0.75 LTV (= minimum required value):	\$2,731,647
- Construction cost:	\$2,140,000
= Supportable site acquisition cost:	\$591,647

of **gross leasable area (GLA)** and an 85 percent **efficiency ratio**, implying a **net rentable area (NRA)** of 29,750 SF. The idea for this project came from a former business partner of Bob's who was interested in leasing a major portion of such a building and believed that Bob was quite capable of delivering on such a development project efficiently and effectively.

Bob believed that, given the proposed design and tenant finishings, the achievable rent in this building would be about \$12/SF in any of several suburban neighborhoods where office buildings of this type were in demand in Midwest City. Preliminary design work indicated that the hard and soft costs, excluding site acquisition, would be about \$2.14 million. Vacancy rates in the types of space markets relevant to this building were running at 8 percent on average. Based on past experience, Bob estimated that operating expenses not passed through to tenants would be \$63,000 per year. For this project, the permanent lender wanted a DCR of 120 percent and an LTV of no more than 75 percent of finished value, with 9.0 percent interest and 20-year amortization. Bob found a potential site for

the project that looked quite promising. The owner of the site wanted \$500,000. Bob needed to decide whether it was worth it. He answered this question according to the steps in Exhibit 17-11, as follows:

As \$591,647 is greater than the \$500,000 price the seller was asking for the site, Bob concluded that this project was financially feasible.

17.4.2 OTHER SHORTCUTS TO EVALUATE A DEVELOPMENT PROJECT

As noted, SFFA is not explicitly in itself a project evaluation methodology. However, there are other common project analysis approaches that do come closer to evaluation. Three such approaches are particularly common, which do allow the computation of land values and related development project profitability: (1) profit margin ratios, (2) enhanced cap rates on cost, and (3) blended long-run IRRs. These are somewhat simplistic rules of thumb that are not as correct and complete as the NPV-based framework we will present in the next chapter. But they are at least quantitative and not unrelated to the project economics.

Consider first the **profit margin** ratio. This is simply the expected gross value of the completed project divided by the total costs of producing the project (land and construction). Present value discounting is usually ignored, and land may or may not be included in the costs in the denominator. If the land is included, it may be included at the historical acquisition cost rather than the current opportunity cost. By this criterion, developers often shoot for something like a 20 percent margin when land is included. That is, they want the expected completed project value to exceed the all-in costs by 20 percent. The intuitive appeal of this measure is obvious, as it seems to present a very direct measure of the “profit” of the development project.³¹

Suppose we expect a property to sell (V_T) for \$10 million after the development is done. The cost (K) to develop is \$6,000,000. Next, say that we assume a 20% margin. How much would be left to purchase the land (L)? Using the profit margin method, we get;

$$1 + \text{Margin} = \frac{V_T}{K + L} = \frac{\$10,000,000}{\$6,000,000 + L} = 1.20$$

$$\text{Land Value (L)} = \frac{V_T}{1 + \text{Margin}} - K = \frac{\$10,000,000}{1.20} - \$6,000,000 = \$2,333,000$$

Thus, applying this method to our project, we find a land value of something like \$2.333 million.

A second rule of thumb developers often apply may be termed **the enhanced cap rate on cost**. As you recall from Chapter 5, a common method of estimating property value is the “direct capitalization” method, in which property NOI (stabilized) is divided by the market “cap rate,” defined as the prevailing current earnings yield based on the NOI observed in transactions in the relevant asset market. For example, in the case of our project described above, assuming an NOI of \$900,000, with a relevant cap rate of 9 percent (found via a “comp” analysis), the property market value of the project is estimated to be \$10 million.

Because of the extra risk in a development project, investors would require a higher cap rate than 9 percent. A typical rule of thumb might require an extra 100 basis points, or in this case, a 10 percent cap rate. However, this cap rate is applied to the total development cost, again, typically without particular care to discounting. Thus, if the developer expects the completed assets to have a market value of \$10 million, with a stabilized NOI of \$900,000, then the project would appear desirable if it

³¹ Sometimes, the profit margin criterion is expressed in a slightly modified form as an “equity multiple” requirement, which is simply the ratio of the expected value of the project benefit (less construction cost or loan OLB) divided by the land and other equity investment value. While no formal present value discounting is implied in the equity multiple measure, the criterion often requires a higher multiple for longer projects.

presented total development costs of no more than $\$900,000/10\% = \$9,000,000$. Construction costs of \$6 million imply that a land price (or net project value at time 0) of no more than \$3 million will make the project attractive.

A third typical ad hoc evaluation procedure is what may be called the **blended long-run IRR**. In this approach, the project's stabilized operating cash flows are projected well beyond the completion of the development phase, typically including a 5 to 10-year horizon beyond stabilization. Thus, we have an overall project cash flow projection over a horizon equal to 5 to 10 years plus the length of the construction phase. A single IRR is then computed over this entire period. The result is, of course, a "blended" IRR that mixes the return earned by the development investment during the development phase with the return earned by the stabilized investment during the operational phase. The latter dominates, both because the capital is fully invested by the time of the operational phase (recall from Chapter 3 that the IRR is a "money-weighted" rate of return) and because the 10-year operational phase projection is typically a much longer period of time than the construction phase. However, the degree to which the stabilized return dominates varies from project to project.

It would not be unusual for developers to require, say, an extra 100 basis points in the long-run blended IRR as compared to a typical otherwise similar stabilized property's going-in IRR (without leverage, or more than that for the equity investment component, assuming projection of permanent debt financing). Thus, for example, in our project, where the market's going-in IRR for the stabilized asset is 9 percent (after applying our tools from Chapter 11), the developer might require a 10 percent going-in IRR on this development project, including a projected 10-year hold after completion. This approach is, therefore, seen to be a rather roundabout way of getting to essentially the same point as the enhanced cap rate on cost approach discussed above. The blended IRR analysis at least requires a more serious look ahead into the future that will determine the project's success or failure. However, it lacks rigor and is no easier to implement (regarding the required information and the number of calculations that must be done) than the rigorous opportunity cost and equilibrium-based procedure we will describe in Chapter 18.

17.4.3 PROBLEMS WITH THE AD HOC APPROACHES

As noted, developers widely use the procedures described above in real-world practice. They are easy to understand with little financial and economic expertise. And we must admit that if applied carefully (and with the benefit of some wisdom borne from experience), such procedures are relevant enough to the actual investment decision economics to have some practical utility, especially as "back of the envelope" tools for preliminary analysis. However, it is essential to recognize how these ad hoc procedures do not provide a *complete or correct financial economic evaluation* of a development project.

The decision to undertake a real estate development project is a capital budgeting decision, in effect, a capital allocation decision at the micro-level, fundamentally like any other micro-level investment decision in a capitalist economy. The **NPV decision rule** and **DCF valuation procedure** that we introduced in Parts I and II of this book are based on the **Wealth Maximization** principle, grounded in the recognition that **Equilibrium** pricing in the relevant asset **Market** provides the relevant **Opportunity Costs**.³²

³² This book is about the *economics* of investment decisions. We do not mean to imply that other perspectives and considerations are not also important, especially for investments that will make a lasting impact on the physical world and have important social and environmental effects. These considerations are rightly more in the domain of public policy and urban planning, and beyond the scope of this text. But economics is also relevant for such other considerations, as private costs of public policies are important considerations, and the efficient use of scarce resources is an important consideration. See Chapter 28 for more explicit consideration of sustainability.

From this perspective, it is easy to see why the developer usually is (or should be) governed by the NPV investment decision rule. Recall from Chapter 4 that this rule requires maximizing the ex-ante NPV over all mutually exclusive feasible alternatives. The basis of the NPV rule in the wealth-maximization principle implies that if the NPV rule is being violated, then money is being left on the table; that is, the developer is reducing their net worth (or shareholder value) as a result of the project development decision, compared to what could have been done instead.

The ad hoc procedures described above fail to account for the NPV rule because they do not quantify the OCC of the development project rigorously and correctly and, therefore, cannot possibly compute the NPV of the project, either explicitly or implicitly. The ad hoc rules of thumb that these procedures represent cannot provide what, in principle, Wall Street wants to know: the expected return to the development project equity to compare this to the relevant opportunity cost of capital (OCC) or even more fundamentally and directly relevant, the OCC-based NPV of the investment.

Of the four procedures described, the blended long-run IRR approach comes closest to a complete analysis. But from an investment analysis perspective, it does not make sense to mix and blend the two very different investment phases of the development project and the stabilized holding. They have very different risk characteristics and so have very different OCC rates. Furthermore, in the modern investment market, typically, different types of investors and/or investment vehicles with different types of investment objectives and risk/return targets provide the source of financial capital for the two different phases, as we described in 17.1. Because of this, the blending of the two investment phases muddy the waters. The developer calculates the long-run blended IRR and then stares at that number, and, if he is honest with himself, really does not know what to make of it. What is the appropriate OCC to which he should compare the IRR he has just calculated? On what rigorous basis would one determine the hurdle IRR by which to evaluate the project? (This is further complicated if financial leverage is employed, and the analysis focuses only on the levered equity.)

To estimate financial feasibility, even the SFFA procedure is needlessly restrictive and complex as a result of the implicit assumption that the completed project will be financed to the maximum possible extent with debt. This may be a plausible enough assumption for small individual developers on Main Street. Still, it is hardly realistic for larger-scale developers, especially those who tap the public equity markets for capital. Because the SFFA procedures do not provide a complete or correct investment evaluation criterion when developers use nothing but such procedures, they are placing the entire burden of making correct capital allocation decisions onto the external capital providers, typically the construction and permanent lenders. Yet these providers of debt capital have no equity directly at risk, and they face little upside potential from a successful project. Traditional lenders are in the business of investing in loans, and loans are much more homogeneous than real estate development equity, making the lending business more efficiently competitive and driving loan NPVs toward zero. It is a fundamental fact that debt investors face different incentives than providers of equity capital do when making optimal capital allocation decisions about physical assets.

17.4.4 How Developers Think about All This

Given this fundamental criticism of the ad hoc methods, we now want to argue (somewhat counterintuitively) that the NPV decision rule must be used *implicitly* among the more successful developers and investors. Furthermore, because of its grounding in market-based principles of opportunity cost and equilibrium, we would argue that actual results in the real world must tend to conform to the implications of the NPV decision rule, at least in places where markets for land, property, and capital function reasonably well. In other words, the simple procedures may not be correct in principle, but maybe somehow, the operation of markets and the desire to maximize wealth tend to lead to reasonably efficient real estate development investment, at least on average, in the big picture.

Why do we think most successful developers use the NPV framework implicitly, if not explicitly? Perhaps we are making a “heroic” assumption. You don’t have to be a genius to see that some pretty bad real estate development decisions get made in the real world. (Of course, hindsight helps, but some projects had to have been turkeys even when the decisions were made.) Nevertheless, we suspect that most real estate development decisions are, in fact, at least approximately consistent with the correct application of the NPV rule (ex-ante), even though the developers and lenders may not be explicitly or consciously thinking of their decision process along the lines of the NPV rule. Why do we say this?

Developers usually face a capital or resource constraint in practice or believe they do. That is, they can only do so many projects at a time. This means they can only do some of the projects that appear to be financially feasible (which they might think they would like to do if they could). As a result, developers are forced, in effect, to rank-order projects and consider feasible combinations of projects. It seems likely that a big part of what makes developers successful is their ability to do this rank-ordering rationally, that is, to choose among mutually exclusive alternatives to pursue those that look best based on the magnitude of the perceived profit or surplus.

In other words, virtually, by definition, successful developers are those who maximize their wealth. And that is what the correct application of the NPV procedure does: by definition and construction, the NPV decision rule maximizes wealth. Thus, successful developers are doing the same thing that the NPV decision rule does. So, implicitly, if not explicitly, they must be applying the NPV rule. Another way of saying this is that applying the NPV rule ensures that developers cover their cost of capital. Less successful developers may not be applying the NPV rule, but they tend to fall by the wayside over time, as any firm that does not cover its cost of capital eventually shrinks or dies.

We would like to suggest that even successful developers may find it useful to make more explicit and quantitative what they have been doing all along implicitly and intuitively, and less successful developers may be able to improve their performance significantly by using the NPV rule. For students who are just starting out in the real estate investment industry, knowledge of the correct framework that underlies wealth-maximizing investment may shortcut the process of learning by trial and error in the real world. Furthermore, as the mainstream investment world of Wall Street and corporate finance grows more relevant to real estate with the rise of REITs and sophisticated private equity funds and the growth of professional, institutional investor interest in real estate, rigorously grounded methods consistent with global standard corporate capital budgeting practice will come to be more widespread and more necessary in the real estate industry.³³

Thus, for the sake of the efficiency of capital allocation in the economy, we turn in the next chapter to a presentation of how to apply fundamental economic principles in a complete model to real estate development investment evaluation and decision-making, including the NPV rule.

17.5 CHAPTER SUMMARY

This chapter builds on the theoretical treatment of land value and development timing presented in Chapter 16 to present an overview of the development decision-making process and the basic methodologies and procedures used in current development project feasibility analysis and construction budgeting. While useful at a practical level, these methods fall short of a rigorously grounded framework for development project investment analysis. This will be our subject in the next chapter.

³³ For a more complete understanding of the market, we should note that there are times when projects that fail to hit minimum NPV are still undertaken. This is a result of the friction costs of laying off a team of professionals during a softer market and then trying to rehire them again once the market improves. This longer run portfolio perspective explains why we see development continuing even when it barely makes sense from a short-term profit perspective.

KEY TERMS

- Site-looking-for-a-use
- Use-looking-for-a-site
- Build to suit
- Lease-up risk
- TI aka build-outs
- Development project phases
- Construction and absorption budget
- Operating budget
- Stabilized (year or cash flows)
- Hard costs
- Soft costs
- Sustainable development
- Absorption or lease-up phase
- On spec (speculative development)
- Tenant build-outs
- Construction loan
- Future advances
- Draw-down (of loan commitment)
- Draw
- Permanent loan
- Certification of occupancy
- Take-out loan
- Accrued interest
- Financing costs (construction loan interest)
- Simple financial feasibility analysis (SFFA)
- Front-door feasibility analysis
- Back-door feasibility analysis
- Gross leasable area (GLA)
- Efficiency ratio
- Net rentable area (NRA)
- Financial feasibility
- Financial economic desirability
- Margin ratio
- Enhanced cap rate on cost
- Blended long-run IRR for developments

STUDY QUESTIONS

Conceptual Questions

- 17.1. Describe the four disciplines or professional perspectives that are involved in the design, analysis, and decision-making process of the typical real estate development project in the private sector.
- 17.2. What are the four phases of development projects? Give a rough indication of the typical relative amount of risk for financial capital invested in each phase by indicating typical expected returns (opportunity cost of capital) at each phase.
- 17.3. What are the major sources of financial capital for each of the development project phases noted above?
- 17.4. Describe the two types of budget that are necessary in the real estate development decision-making process.

- 17.5. Differentiate between hard and soft development costs.
- 17.6. How does a *construction* loan differ from a *permanent* mortgage loan? What is a *mini-perm* loan?
- 17.7. Describe the *front-door approach* to simple financial feasibility analysis. How does this differ from the *back-door approach*?
- 17.8. We argue in this chapter that SFFA analyses serve as useful preliminary financial analysis tools but should not be relied upon as the ultimate or complete model of the economic desirability or profitability of a proposed development project. What are the major problems or limitations of the simple financial feasibility analysis (SFFA) techniques in this regard?

Quantitative Problems

- 17.9. You wish to build an office/warehouse project, also known as R&D space or flex space. Market rents seem to be around \$15/SF/year with a 5 percent vacancy rate in the local area. All expenses are passed through to tenants except property taxes, insurance, and management, which you estimate at \$5/SF/year. Mortgage rates are 11 percent for a 20-year loan with a five-year balloon. Construction costs for your planned 20,000-gross-leaseable-square-foot project are estimated at \$1,030,075 in total. All 20,000 SF are rentable. The debt service coverage ratio required is 120 percent, and the maximum LTV ratio is 75 percent.

Use the SFFA *back-door* procedure to determine what you could pay for the land.

Now assume total construction and site acquisition costs (i.e., total development costs) are \$1.3 million. Use the *front-door* SFFA procedure to determine what the required rents would be per square foot. Recalculate the required rent assuming a vacancy rate of 15 percent.

- 17.10. The following table shows the projected draws to pay the construction costs of a project that is expected to take four months to complete. The draws are projected to occur at the beginning of each month. What is the projected balance due on a construction loan at the end of the fourth month if the interest rate is 10 percent per annum with monthly compounding?

Month	Draw
1	\$1,000,000
2	\$2,000,000
3	\$1,500,000
4	\$2,500,000

18 Investment Analysis of Real Estate Development Projects, Part 2

Economic Analysis

CHAPTER OUTLINE

- 18.1 The Basic Idea
 - 18.1.1 Applying NPV to Development Projects
 - 18.1.2 Operational Leverage and Estimation of the OCC for Development Investments
 - 18.1.3 The Development Risk Ratio and Implications for Characteristic Development Regions
- *18.2 Advanced Topics in Development
 - *18.2.1 The Relationship of Development Valuation to the Real Option Model of Land Value
 - *18.2.2 Capital Structure for Real Estate Development
 - *18.2.3 “Soup to Nuts”: Including Lease-up and Land Assembly: “*Unblending*” the Blended IRR
- 18.3 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The more rigorously correct NPV-based approach to financial evaluation of development projects.
- The relationship of operational leverage to development project risk and return.
- How to estimate the opportunity cost of capital or reasonable expected return for real estate development project investments.

The previous chapter introduced you to some of the major tools commonly used to examine the financial feasibility of real estate development projects. We also suggested, however, that those tools need to be improved in some important respects. Financial “feasibility” is not exactly like “desirability.” The tools used in current practice tend to be *ad hoc*, not grounded solidly in a rigorous economic framework, and generally not fully considering the basic economic principles of market equilibrium and opportunity cost. This not only makes the *ad hoc* procedures subject to possible error or incompleteness, but it makes it difficult for real estate development investment analysis to communicate and relate coherently to the mainstream investments and corporate finance world, which is the source of ever more of the capital used to finance development (including not just debt but in recent years much of the equity capital). Finally, without a rigorous economic framework, the analyst and the decision maker lack the kind of deep understanding that can facilitate greater creativity and innovation in project conceptualization and deal formation.

It is the job of this chapter to remedy this deficiency. In earlier chapters, we noted that the investment analysis of real estate development projects is essentially an exercise in capital budgeting, a term

borrowed from corporate finance literature. Capital budgeting is the field from which the NPV rule arose. It is a micro-level investment decision arena whose basic concepts and principles were treated in Part III of this book, particularly in Chapters 14 and 15. On top of that, Chapter 16 presented some important fundamental economic principles that relate particularly to real estate development project evaluation and decision-making. What remains for us to do in the present chapter is to extend our treatment of this topic to a more practical level. In particular, we will examine how to apply the NPV investment decision procedure to the typical commercial property development project. We will do this in a framework that is consistent with the real options model we introduced in Chapter 16.

18.1 THE BASIC IDEA

As we have said, the essential framework for applying sound financial economic theory to the investment evaluation of development projects is correctly applying the NPV rule. The basic NPV rule was presented and discussed in Chapter 4. In Chapters 14 and 15, we extended the basic rule to encompass the use of debt financing, using the value additivity principle and the APV (adjusted present value) rule as a useful shortcut that can avoid unnecessary muddying of the waters when debt is employed. The principles and procedures introduced in those two chapters completely apply to the financial evaluation of development projects. In addition, our discussions of leverage in Chapter 12 and of the real option model of land value and optimal development timing in Chapter 16 are particularly relevant to development investment evaluation. This section will get a bit more specific about this application. In particular, development projects are typically characterized by three somewhat unique features as compared to investments in existing, fully operational properties:

1. **Time-to-build:** In development projects, the investment cash outflow is spread out in time instead of occurring all at once upfront. This gives development investments inherent “operational leverage,” even if no financial leverage is employed.
2. **Construction loans:** Debt financing is almost universal in the construction phase of a typical development project, and this debt typically covers all of the construction costs.
3. **Phased risk regimes:** Because of the operational leverage noted above, development investment typically involves very different levels of investment risk between the construction (or development) phase, the absorption (or lease-up) phase, and the long-term (stabilized or permanent) phase when the completed project is fully operational. The result, as described at a broad-brush level in the first section of Chapter 17, is a very different economic opportunity cost of capital (OCC) and typically different sources of financial capital for the different phases.

The correct application of the NPV rule to development projects requires a consideration of these three unique features. (See 10-minute Video 1 in Chapter 18 Supplements folder.)

18.1.1 APPLYING NPV TO DEVELOPMENT PROJECTS

To begin this consideration, recall that in Chapter 4, we defined the NPV of the investment from the buyer's perspective as being equal to the difference between the value of the property being obtained and the cost of obtaining it, which we labeled as

Equation (1)

$$NPV = V - P$$

In the application to investment in existing properties, V is the value of the operating property being obtained, and P is the price of obtaining the property.

In order for the NPV rule to make sense, both the benefit and the cost must be measured in present, certainty-equivalent dollars, that is, in dollars adjusted for risk as of the time the investment

decision is being made. Otherwise, we would be comparing apples to oranges. By labeling the present (decision time) as time 0 , we can clarify the NPV evaluation by subscripting the benefit and cost measures to indicate the point when they are measured, as follows:

Equation 1a

$$NPV_0 = V_0 - P_0$$

In the case of investments in fully operational (stabilized) properties, it is relatively easy to estimate both the benefit and cost as of time 0 and thereby quantify equation (1a) directly. In the case of development projects, quantifying (1a) is a bit more complicated because neither the benefit nor all of the cost is incurred upfront at time 0 .

To see this point, let's consider a simple example. Hereandnow Place is a newly completed office property with two buildings and a joint parking lot. Its expected property-before-tax cash flow (PBTCF) is \$75,000 per month (\$900,000/year, paid monthly) in perpetuity. The property investment market's required expected total return for assets like Hereandnow Place is 9 percent. In other words, the opportunity cost of capital—OCC—for investment in Hereandnow Place is 9 percent per year. We will assume that all rates of return are quoted as nominal annual rates with monthly compounding. Therefore, the 9 percent OCC equals $9\%/12 = 0.75\%$ per month simple return. Thus, Hereandnow Place is worth \$10 million today:

$$\$10,000,000 = \frac{\$75,000}{1.0075} + \frac{\$75,000}{1.0075^2} + \dots = \frac{\$75,000}{0.0075}$$

Because markets are competitive, it is likely that Hereandnow Place would sell for \$10 million, and an investment in Hereandnow Place would therefore be made at zero NPV (at least from a market value perspective). Thus, the NPV of such an investment would be easily quantified as:

Equation 1b

$$NPV_0 = V_0 - P_0 = \$10,000,000 - \$10,000,000 = 0$$

At the \$10 million price, Hereandnow Place provides the investor with an expected return (going-in IRR) of 0.75 percent per month, or 9 percent per year (ENAR), or $1.0075^{12} - 1 = 9.38\%$ effective annual rate (EAR). Assuming that the risk-free interest rate in the capital market is 3 percent (ENAR, or 3.04 percent EAR), this implies a market risk premium in such an investment in stabilized property equal to 6 percent (nominal, or $6.34\% = 9.38\% - 3.04\%$ EAR).

Now, across the street from Hereandnow Place is a vacant and developable land parcel identical to the site on which Hereandnow Place is built. The market for such office properties at this location is still strong, and it is clear that to build an essentially identical development to Hereandnow Place on this new site would indeed be the highest and best use of the site and that it would make sense to commit irreversibly to such a development right away. Indeed, to eliminate lease-up risk as a source of difference between the development versus stabilized asset, we will assume that the new development is already preleased (that is, a forward leasing commitment has been signed) with rental terms identical to those in Hereandnow Place. The new development, to be called FutureSpace Center, would involve construction costs that are contractually fixed and would have to be paid at four points in time, quarterly payments of \$1.5 million each at the end of 3, 6, 9, and 12 months from the time when the development decision would be made ("time 0 ," which we assume is now). The result would be that the first of the twin buildings would be completed and ready for occupancy six months from time 0 , and the second building would be completed and ready for occupancy 12 months from time 0 . Each FutureSpace Center building, which would be effectively just like the Hereandnow buildings, is expected to produce a net cash flow of \$37,500 per month (\$450,000/year) each from its time of completion, in perpetuity, based on the already committed leasing arrangements.

The essential development-investment-valuation question in this context is: *What is the price that can be paid today for the FutureSpace land site such that the development investment will be zero NPV?* This is the value of the land, the price the FutureSpace land site would sell for in a competitive market. This question may be equivalently framed as: *What is the NPV of the development project investment apart from the land cost?* The answer to the question phrased in this way, tells us not only what the maximum price is that can be paid for the land to produce a nonnegative NPV, but also enables us to quantify what the NPV of the development investment would be given any specified land price.

To answer this question, we need to quantify V_0 and P_0 for the FutureSpace development project, the certainty-equivalent values of the benefits and costs of the project as of time 0, the present time. The benefits are, of course, the two new office buildings that will be built, and the costs, apart from the land, are the construction costs required to build the two buildings. We will label those construction costs as K . All we have to do is forecast the magnitudes and timings of these benefits and costs and then discount those projected values back to the present, using OCC discount rates appropriate to each type of projected cash flow.

Consider first the benefits of the to-be-completed office buildings similar to Hereandnow Place. The first FutureSpace building (named “FutureSpace One”) will be completed in six months. At that time, it will represent an asset projected to be worth \$5 million. This is because FutureSpace One will have the same OCC as Hereandnow Place (9 percent per annum nominal, or 0.75 percent per month simple) because FutureSpace One will be essentially the same type of asset as Hereandnow Place and, therefore, will present investors with the same amount of investment risk. Thus, upon its completion, FutureSpace One is expected to be worth:

$$\$5,000,000 = \frac{\$37,500}{1.0075} + \frac{\$37,500}{1.0075^2} + \dots = \frac{\$37,500}{0.0075}$$

Similarly, the second new building, FutureSpace Two, is expected to be worth \$5 million upon completion in 12 months. Thus, we can compute V_0 , the present value of the benefits of the development project, as \$9.352 million, as follows:¹

$$V_0 = \$9,352,000 = \frac{\$5,000,000}{1.0075^6} + \frac{\$5,000,000}{1.0075^{12}}$$

Note that we are simply projecting the values and timing of completion of the assets being produced in the development project and discounting those values back to present at the OCC rate appropriate to investments in such assets. It is as if we were actually planning to sell each building as soon as it is completed. In fact, we could do that, though we need not feel pre-committed to doing so. We would get the same answer if we projected receiving the net rents from the first building for the six months from months 7 through 12, at which time we would have two assets worth \$10 million together:

$$V_0 = \$9,352,000 = \sum_{t=1}^{12} \frac{\$37,500}{1.0075^t} + \frac{\$10,000,000}{1.0075^{12}}$$

Expressed this way, the equation suggests that we plan to sell the entire FutureSpace Center project upon its completion in 12 months. In fact, however, this is not necessarily implied by the equation. Whether we sell or keep FutureSpace Center, in month 12, the development project is expected to produce an asset worth \$10 million (as well as six months’ worth of net rents by then). We can decide later what to do with the asset.

¹ Note that what we are referring to as the gross “benefits” of the development project here equates to what we called in Chapter 16, the “underlying asset” of the land development *option*.

The key point here at time 0 is that month 12 is the end of the investment's development (or construction) phase. From that time onward, any continued investment on our part (or any new investment by a purchaser) would be a standard investment in a stabilized operating property as described in equation (1b), like what we are familiar with analyzing in Chapters 4 and 11. Here, we are interested in analyzing development project investments, so cutting off the analysis horizon after month 12 makes sense. As described previously, this is important from an investment-analysis perspective because the development phase investment has very different risk and return characteristics than the stabilized operating phase investment. Different types of investors will typically be interested in these two different phases, and “apples-to-apples” comparisons of risks and returns require that we not mix investments of very different risk-and-return characteristics into a single, undifferentiated analysis frame.

Note that the present value of the gross benefits, V_0 , is less for the development project FutureSpace Center than for the stabilized property Hereandnow Place across the street (\$9.352 million versus \$10 million), even though the buildings are identical. This is because part of Hereandnow Place's value is the present value of 12 months' worth of \$75,000/month net cash flow from months 1 through 12. During the 12-month development phase, the FutureSpace Center project will only generate half that rate of monthly cash flow, and for only the last six of the 12 months starting from time 0 . Obviously, this is because the FutureSpace buildings need to be built, whereas the Hereandnow buildings are already complete.

Now, let's consider the cost side of the NPV equation. The FutureSpace Center project will require \$6,000,000 in construction costs, projected as four equal payments of \$1.5 million at the end of months 3, 6, 9, and 12. Given that these expenditures will happen in the future, we must also discount these costs. Remember from Chapter 4 that the discount rate (OCC) must reflect the amount of risk in the asset values or cash flows being discounted, *as the capital market evaluates that risk*. In this context, “risk” would describe the uncertainty of the total construction bill the investor will ultimately face. (Which in turn would be a threat to its profits.) The obvious follow-up question is: at what rate should we discount construction costs? Higher or lower compared to the project's OCC of 9%? A first knee-jerk reaction might be *higher* because, as we all know, developments are risky! But in fact, the OCC of construction costs is actually lower (almost always much lower) compared to the stabilized property's OCC because they are more stable, predictable, and idiosyncratic. More specifically;

- **Construction costs are relatively stable.** Construction costs are less uncertain in the future compared to real estate values. In most countries, the “price elasticity” of the supply of construction goods and services (labor and materials) is quite high. This means that a small price change will evoke a large change in the quantity supplied. This stabilizes prices. The underlying reason for this is that the construction market is *very competitive*. Due to this competitiveness (and without collusion), no single builder can increase its prices above the cost of producing the building, even if demand for their services increases. When demand falls, construction firms will lay off workers, stop ordering materials, or even drop out of the business altogether. As a result, construction costs typically follow inflation, and their cyclical is much attenuated. The opposite is true with built property, as land is very inelastic (due to supply constraints; see Chapter 7). As a result, real estate asset values tend to be cyclical and more volatile.
- **Better risk mitigation.** The risk that does exist (i.e., cost overruns) tends to be *within* a given project. (This is referred to as “specific” or “idiosyncratic” risk, as distinct from “systematic” risk.) Surprises in construction costs are likely more related to (unforeseen) technical and physical issues than financial issues. Because these risks are idiosyncratic with little *covariance* with financial market returns in the capital markets, investors can easily diversify away from these risks.² Furthermore, the contractor may provide a guarantee against such technical risks, relieving the developer of such risks altogether. Even cost overruns caused by changes in

² In Chapter 23 on the subject of equilibrium asset pricing (found in the online portion of this book), we will see that the capital market tends to only pay for risk (in the form of higher expected returns) when the risk cannot be diversified away.

material prices (which typically are, but not necessarily, owed by the developer) can be hedged away by investing in futures of material prices.

For these reasons, the OCC for construction costs is *closer to the risk-free* rate than the OCC for the stabilized asset being built. This creates what we called in Chapter 12 *positive leverage*. As we saw in Chapter 12 (and based on fundamental equilibrium theory), positive leverage implies not just greater expected return (OCC) but also greater investment risk (and vice versa in equilibrium). Given that construction costs are typically a large fraction of the total value, developments tend to have high leverage and, therefore, high risk. To provide more intuition on why using a lower construction cost OCC does not contradict the statement that development is riskier, consider using a high OCC for construction costs. In that case, you will find:

- A lower present value magnitude at Time 0 for future construction costs paid at and future Time T;
- A smaller PV construction cost amount, therefore, to subtract from the PV of the project benefits in calculating the NPV of the project as of Time 0; thereby
- Making the project appear *more* profitable (higher NPV than if we had used a lower OCC rate to discount the construction costs), which is
- Just the opposite of what we want to do to reflect the risk of construction!

Even if we haven't convinced you that the low OCC for construction costs does not mean there is little risk in development (and we adore the skeptical student), understand that we have been in this situation before without explicitly mentioning it. In Chapters 12 and 13, we discussed the impact of financial leverage on real estate investment. We also discussed how the debt service payments (from a lender's perspective) are commonly less risky than property cash flows, as banks have a senior claim on the cash flows, and the payments are fixed upfront. However, at the same time, we discussed how leverage increases the risk and return of the property investment. Thus, if you believed us there, why not here?³

With this OCC question out of the way, let's assume a risk-free interest rate of 3 percent per annum nominal (with monthly compounding, or $3\%/12 = 0.25\%$ per month simple). We can now easily quantify the time 0 value of the construction cost component of our NPV equation:

$$K_0 = \$5,889,000 = \frac{\$1,500,000}{1.0025^3} + \frac{\$1,500,000}{1.0025^6} + \frac{\$1,500,000}{1.0025^9} + \frac{\$1,500,000}{1.0025^{12}}$$

Thus, the NPV of the development project investment as of time 0, apart from land cost, is:

Equation 2

$$V_0 - K_0 = \$9,352,000 - \$5,889,000 = \$3,463,000$$

This value of \$3.463 million should be interpreted as the net economic value of the particular development program represented by the FutureSpace Center project, as we have specified that project.

³ All of this is nothing more or less than equilibrium in the capital market: the Security Market Line introduced in Chapter 12: more risk ==> more expected return, and vice versa. (aka, "You can't get something for nothing!") But perhaps we convinced you so much that now you are wondering why we would have to discount the construction costs at all. Say that we know (for sure) that we must pay \$10 million to the contractor in 1 year. If the risk-free rate is 3 percent, that would mean a PV of \$9.7 million (rounded as \$10m/(1.03)). In this case, you could (re)invest the \$9.7 million now in short-term T-bills for a year, which would be worth \$10 million in a year. In other words, the OCC includes two components, the risk premium and also the risk-free rate (time value of money): $r = rf + RP$.

Suppose the price of the FutureSpace Center land site is \$3.463 million. In that case, the total economic cost of the investment in the development project, P_0 , will be \$9.352 million, and the investment will have a zero NPV as of today, the time when the development investment decision must be made. If the price of the land is any value x , then the NPV_0 will be $\$3,463,000 - x$. If this is nonnegative, then it makes *economic* sense to undertake the development investment, given that it is the highest and best use of the site.⁴

TEXT BOX 18.1 A BRIEF DIGRESSION: HIGHEST AND BEST USE ANALYSIS...

Although our primary focus in this section is on investment in the land commitment and construction phase of development projects, it may be worthwhile to step back for a brief digression to consider the applicability of this analysis framework to the preliminary phase. Recall that we specified at the outset of our description of the FutureSpace Center example that this project is the highest and best use (HBU) of the site. Referring back to the multistage overall development process outlined in Section 17.1 of the previous chapter, the determination that FutureSpace is indeed the HBU for this site would have been made during the preliminary phase, just prior to what we are calling *time 0*. During the preliminary phase, explicit HBU analysis should be based on the same metric as we quantified in equation (2), the net economic value of the project (exclusive of land cost), which we labeled " NPV_0 " in that context. The HBU for the site is the project design that maximizes this NPV_0 exclusive of land cost.

Normally, the value of the land, which we have labeled " x " above, would reflect that HBU. (Note that this is consistent with the optimal density model we presented in the text box in Section 17.1.1. This implies a zero NPV for the HBU and a negative NPV for any other use, including the land cost in the NPV. If the land value exceeds the $V_0 - K_0$ value for the HBU, then the site is not yet "ripe" for development; we are still in the land speculation real option phase.) However, an entrepreneurial developer can envision, design, and bring to site feasibility (e.g., by land assembly and permitting) a more valuable project (i.e., that has a higher equation (2) NPV_0 exclusive of land cost) than anyone else could do. Through such a process, the developer adds value during the preliminary phase of the overall development process. This added value is ultimately reflected in the *investment* value of the land (as assembled and permitted) for the particular developer. However, as introduced in Chapter 14 (see 14.1.2), it may not be entirely reflected in the *market-value-based* opportunity cost of the land, x , if no other developer could execute the project. In any case, the fundamental NPV decision rule of maximizing the NPV over all mutually exclusive alternatives governs the HBU analysis in the preliminary phase of the overall development process, with NPV defined as we have done here, exclusive of land cost.

When a project makes economic sense, it should be possible to find some combination of participants and structure the financing and control of the deal in a manner that makes the project

⁴ Recall from Chapter 16 that consideration of option value could make it optimal to wait before investing in the development project. Such "option premium" value and optimal waiting should normally be reflected in the market value of the land. Hence, it would be reflected in a value of x greater than \$3.463 million in this case. However, recall that at the outset of our discussion of this example, we posited that it was indeed optimal to commit irreversibly to immediate development of the FutureSpace Center project. To relate this to the discussion in Chapter 16, the current underlying asset value of \$9.352 million (signaled by an observable comparable completed asset current value of \$10 million as seen in the value of Hereandnow Place) apparently equals at least the "hurdle value" that makes immediate development optimal, as described in section 16.5.1.

feasible and provides fair returns to all parties ex-ante. When a project does not make economic sense, it will not be possible to assemble such a package to make the project feasible without at least one party (one contributor) not achieving a fair expected return going into the deal.

18.1.2 OPERATIONAL LEVERAGE AND ESTIMATION OF THE OCC FOR DEVELOPMENT INVESTMENTS

(See 4-minute Video 2 in Chapter 18 Supplements folder.)

We have previously noted that a salient feature of development project investments is that, because of the operational leverage inherent in the project, the risk and return characteristics of such investments are very different from those of stabilized operating properties. In particular, development phase investments are riskier and, therefore, must provide higher expected returns than investments in stabilized properties. Understanding this difference is very important because (as we have repeatedly noted in this book) the supply side of the capital market (those who have money to invest) is divided into different types of investors who are looking for different levels of ex-ante risk and return in their investments. This section will show you how to rigorously quantify the difference between development phase investment risk and stabilized operating phase investment risk based on fundamental economic principles.

The extra risk in development investments is easily seen in our previous FutureSpace Center project illustration. We already ascertained that \$3.463 million is the zero-NPV fair market value of the upfront investment in the FutureSpace Center development project (the opportunity value of the land). We can, therefore, easily compute the fair market expected return on the development phase investment in FutureSpace Center. The going-in expected return on this development investment is seen to be 16.59 percent per annum EAR (1.29%/month simple, or 15.44% ENAR with monthly compounding) by computing the IRR of the development phase cash flows as follows:⁵

$$\$3,463,000 = \frac{\$1,500,000}{(1+IRR/mo)^3} + \frac{\$1,500,000}{(1+IRR/mo)^6} + \frac{\$1,500,000}{(1+IRR/mo)^9} + \frac{\$1,500,000}{(1+IRR/mo)^{12}}$$

$$\Rightarrow IRR = 1.287\%/mo, \Rightarrow (1.01287)^{12} - 1 = 16.59\%/yr$$

Now compare this to the expected return on an investment in the otherwise identical stabilized property across the street at Hereandnow Place. Recall that the zero-NPV (fair market value) expected return on investment in Hereandnow Place is 9 percent (or 9.38 percent EAR). Considering the 3.04 percent risk-free interest rate, the expected return risk premia for the two investments are (in EAR terms):

$$9.38\% - 3.04\% = 6.34\%, \text{ for Hereandnow;}$$

$$16.59\% - 3.04\% = 13.55\%, \text{ for FutureSpace}$$

Thus, the development project investment has a risk premium that is $13.55/6.34 = 2.14$ times that of the stabilized property. Given that both of these are fair market return risk premia and, therefore, must provide the same expected return risk premium *per unit of risk*, the implication is that the

⁵ Here we are assuming sale (or valuation realization) of each building upon completion. Thus, the net cash flow in months 6 and 12 is $\$5,000,000 - \$1,500,000 = \$3,500,000$. Other approaches to realization of the project benefits will yield different IRRs. A “canonical” approach to quantifying consistently development project investment returns will be discussed later in this section.

development investment has 2.14 times the risk of the stabilized property investment (as the capital market perceives and evaluates risk).

This difference reflects the **operational leverage** in the development project. It exists even in development projects that are entirely preleased, such that there is no lease-up or rental market risk. Essentially, operational leverage occurs whenever an investment involves future cost outlays that are not perfectly correlated with the values of the future gross benefits to be obtained from the investment. In other words (using our previous NPV terminology), operational leverage arises whenever P does not occur entirely at time 0 and is not perfectly and positively correlated with the subsequent realization of V . In general, the greater the relative magnitude of the construction costs, K , compared to V , and the earlier the realization of asset value, V , occurs in time relative to the incurring of K , the greater the amount of operational leverage in the development project. For example, other things equal, a developer can increase operational leverage by selling off completed assets sooner. Investments in stabilized properties have no operational leverage because the investment cost occurs entirely at time 0 , the same time as asset value is realized.⁶

To make this point more concrete, let's suppose that in our FutureSpace and Hereandnow examples, the ex-post asset values turn out to be 10 percent less than had been expected at time 0 . Instead of achieving its expected IRR of 9.38 percent, the investment in Hereandnow would actually achieve a lower return, namely -1.04 percent, computed as:

$$\begin{aligned} \$10,000,000 &= \sum_{t=1}^{12} \frac{\$75,000}{(1 + IRR/mo)^t} + \frac{\$9,000,000}{(1 + IRR/mo)^{12}} \\ \Rightarrow IRR &= -0.087\%/mo, \Rightarrow (1.99913)^{12} - 1 = -1.04\%/yr \end{aligned}$$

Thus, the ex-post investment result for the year would fall short of the ex-ante expectation by $9.38\% - (-1.04\%) = 10.42\%$. But the investment in FutureSpace Center would suffer much more on a percentage basis, experiencing an actual return of -13.42 percent in the ex-post IRR calculation based on valuations of the two buildings at $\$4.5$ million instead of $\$5$ million at the time of completion in months 6 and 12:

$$\begin{aligned} \$3,463,000 &= \frac{-\$1,500,000}{(1 + IRR/mo)^3} + \frac{\$3,000,000}{(1 + IRR/mo)^6} + \frac{-\$1,500,000}{(1 + IRR/mo)^9} + \frac{\$3,000,000}{(1 + IRR/mo)^{12}} \\ \Rightarrow IRR &= -1.19\%/mo, \Rightarrow (0.9881)^{12} - 1 = -13.42\%/yr \end{aligned}$$

Compared to FutureSpace's ex-ante expected return of 16.59 percent, this is a 30.01 percent return shortfall, some 2.9 times the 10.42 percent return shortfall of the investment in the stabilized property. In other words, the particular type of risk event described here is "levered"

⁶ Obviously, the investor can add *financial leverage* by the use of debt financing of the investment, as described in Chapter 12. But here, to illustrate fundamental principles, we are considering investment without the use of debt finance. Also, some properties may have a type of operational leverage in their stabilized phase if the landlord is responsible for paying operating expenses that are relatively fixed while the rental income is variable and not correlated with those expenses. One example is a property subject to ground lease rental payments (as described in the supplemental materials for Chapter 13). However, to the extent this type of stabilized phase operational leverage exists, it should be completely reflected in the stabilized phase OCC (e.g., the 9.38% in our previous example) and, therefore, fully accounted for in the value of the stabilized asset, V .

almost three times in the development investment return compared to the stabilized property investment return.⁷

This kind of effect is essentially the same type of impact on risk and return as was described in our discussion of leverage in Chapter 12. The effect described in Chapter 12, however, was due to *financial leverage* (use of debt financing of the investment), whereas here no debt financing is being employed (hence, no financial leverage).

With this in mind, let us return to our question of how to compute the OCC of development-phase investments. It should now be clear that the development-phase OCC is greater than that in an otherwise similar stabilized property because the investment risk is greater due to the operational leverage. But how much greater? If we know the OCC for the stabilized property being built, how can we compute the OCC for an investment in the development phase?

Before we answer this question, let's take a minute and put it in perspective. Why do we feel a need even to ask this question? In investment analysis, we generally want to know the OCC of an investment in order to use the OCC as the discount rate for purposes of computing the NPV of the investment. However, as our FutureSpace project illustration shows, we do not need to know the development phase OCC itself in order to compute the NPV of the development phase investment. We can compute the development NPV by separately discounting the gross benefit (V) and cost (K) cash flow (and/or asset value) streams, using OCC rates appropriate to each, as described in Section 18.1.1. In this context, the development phase OCC emerges as an *output* from the NPV computation rather than an *input*. For example, we determined that the OCC of the FutureSpace Center development project was 16.59 percent.

Yet this 16.59 percent output was seen to be an interesting number. We used it to quantify the amount of investment risk in the FutureSpace project relative to that in an otherwise similar stabilized property investment, based on the principle that the expected return risk premium must be the same *per unit of risk* across these different types of investments. This principle applies not just from a normative perspective but fundamentally in a positive (empirical) sense because investors can trade across and between the markets for these different types of investments. That is, they can buy and sell development investments and investments in stabilized properties. If the two types of investments do not present the same going-in expected return risk premia per unit of risk (the same **price of risk**), then investors will tend to sell the type of investment that presents the lower risk premium per unit of risk and buy the type of investment that presents the higher risk premium per unit of risk. Over time, this will drive market prices of the different types of investments to levels that will present investors with the same risk premium per unit of risk.

If we agree that development phase investment OCCs are interesting numbers and that we know how to derive them as *outputs* from the NPV analysis described above, we still face a practical consideration that can be awkward for some analytical purposes. In particular, the specific numerical value of the development phase OCC depends on the particular cash flow and value realization timing assumptions employed in the IRR computation. For example, the 16.59 percent OCC we derived for the FutureSpace project was based on the cash flow timing assumptions of equation (3a). These included the assumptions that:

- Each building's complete value would be fully realized (including the present value of any subsequent net rents) when, and only when, each building was completed (in months 6 and 12) and:

⁷ This ratio is greater than the 2.14 ratio in the ex-ante risk premia computed earlier, because here we are making a bit of an “apples-vs.-oranges” comparison. We are hypothesizing that the Hereandnow investment does not take the “hit” until 12 months from now in its asset value at that time, while the FutureSpace investment takes half of the hit (in the FutureSpace One building value) upon its completion in only six months and the other half in 12 months. The earlier registration of the loss causes a larger impact on the return.

- The construction costs would be paid during the project as they were incurred (months 3, 6, 9, and 12).

We could have made alternative assumptions. For example, suppose we replaced the first assumption above with an assumption that asset values would be realized only at the end of the complete two-building development phase at month 12 (for both buildings). Then our going-in IRR calculation would be as follows:

$$\begin{aligned} \$3,463,000 &= \frac{-\$1,500,000}{(1+IRR/mo)^3} + \frac{-\$1,500,000}{(1+IRR/mo)^6} + \sum_{t=7}^8 \frac{\$37,500}{(1+IRR/mo)^t} + \frac{-\$1,463,000}{(1+IRR/mo)^9} \\ &\quad + \sum_{t=10}^{11} \frac{\$37,500}{(1+IRR/mo)^t} + \frac{\$8,537,500}{(1+IRR/mo)^{12}} \end{aligned}$$

$$\Rightarrow IRR = 1.07\%/mo, \Rightarrow (1.0107)^{12} - 1 = 13.58\%/yr$$

The resulting IRR of 13.58 percent is quite different from our previous calculation of 16.59 percent, yet it is clearly a valid representation of the expected return on the development investment. Which IRR represents the “true OCC” of the development phase investment?

In fact, they can both be true. As long as they represent feasible and plausible strategies for the actual cash flow realizations of the development phase investment, any number of valid OCC rates can be generated for a typical development phase investment. Comparisons among these rates can provide some insight as to how to manage and finance the project so as to achieve a target risk/return profile. For example, as we have just seen, planning to sell each building upon its completion will add to both the expected return and the risk in the FutureSpace project, as indicated by the difference between the 16.59 percent versus 13.58 percent IRRs computed in equations (3a) and (3b).

Nevertheless, it can be useful to have a more standardized formula for defining the OCC of development project investments. This can provide a common framework and, in some contexts, facilitate comparisons across projects. Fortunately, an approach to standardizing the measurement of development projects, such as OCC, naturally suggests itself in the very widespread use of construction loans to finance the construction cost component of development project investment. As noted in Chapter 17, the classical construction loan results in all of the construction cost cash outflow being paid by the developer (the loan borrower) at one point, namely, upon completion of construction (when the construction loan is due). This dovetails nicely with the intuitively appealing conceptualization of the net benefits of the development being realized as of a single future point in time, namely, the time of completion of construction of the entire project.

Thus, a sort of “canonical” development cash flow pattern can be proposed, in which all cash flow occurs at two, and only two, points in time: (i) “time 0” (the present) when the irreversible commitment to the development project must be made, and the cost of the land is effectively incurred;⁸ And (ii) “time T” when the construction is essentially completed, and the developer obtains the net difference between the gross value of the built property as of time T minus the construction costs compounded up to time T.⁹ This canonical assumption of land investment at time 0 and net construction profits realized all at time T is not particularly helpful for complex, multistage projects or for projects that involve the production and sale of many separate assets over a long

⁸ Remember that the economic opportunity cost of the land at time 0 is its market value as of that time (what it could be sold for as a developable land parcel), not necessarily equal to its historical acquisition cost.

⁹ The time between time 0 and time T, the duration of construction, is what we labeled in section 16A.1.3 (in the supplement materials to Chapter 16) as the “time-to-build” (ttb).

period of time (such as most single-family housing developments). However, it provides an intuitively appealing picture of many development projects.

The **canonical formula** for the OCC of development investments can be expressed by the following condition of equilibrium across the markets for developable land, built property, and contractually fixed cash flows or debt assets (see Exhibit 18-1):

Equation 4

$$\frac{V_T - K_T}{(1+E[r_C])^T} = \frac{V_T}{(1+E[r_V])^T} - \frac{K_T}{(1+E[r_D])^T}$$

where:

V_T = Gross value of the completed building(s) as of time T .

K_T = Total construction costs compounded to time T .

$E[r_V]$ = OCC of the completed building(s).

$E[r_D]$ = OCC of the construction costs.

$E[r_C]$ = OCC of the development phase investment.

T = The time required for construction.

In most cases, all of these variables can be observed or estimated with relatively high confidence, except for the OCC of the development phase investment, $E[r_C]$. Solving (4) for $E[r_C]$ we obtain:

Equation 4a

$$E[r_C] = \left[\frac{(V_T - K_T)(1+E[r_V])^T (1+E[r_D])^T}{(1+E[r_D])^T V_T - (1+E[r_V])^T K_T} \right]^{\frac{1}{T}} - 1$$

Another way of writing this is to define $NPV_0 = V_0 - K_0$ as the result of equation (4), the net economic value of the project exclusive of land cost, which can be computed from either the right-hand or the left-hand side of equation (4). Computing it from the right-hand side, we can then solve for the development OCC as:

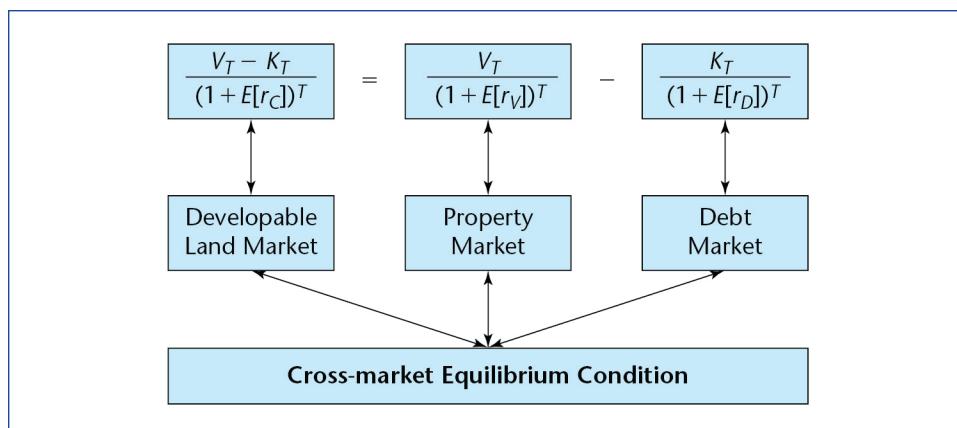


EXHIBIT 18-1 The “Canonical” Formula.

Equation 4b

$$1 + E[r_c] = \left(\frac{V_T - K_T}{V_0 - K_0} \right)^{\left(\frac{1}{T} \right)}$$

In equation (4), the left-hand side represents the investment in developable land, which produces the net difference between the built property value and its construction cost at time T . The right-hand side of equation (4) represents essentially a way to duplicate this development investment by investing in a combination of the following: a long position in built property of the type being developed, and a short position (borrowing) in an asset that pays contractually fixed cash flows (debt) in the amount of the construction costs of the project. The equation must hold for investments across these three types of asset markets (land, built property, and debt) to provide investors with competitive risk-adjusted return expectations (that is, equal expected-return-risk premia per unit of risk across the three types of investments).

Equation (4) represents the same argument and model as in Section 16.3, where we presented the real option model of land value and optimal development timing. The only difference is that here we assume that the timing is ripe to begin the development project, so the development investment OCC is based on a definite commitment to proceed with the project.¹⁰

To see how equation (4) works in practice, let us apply it to our previous illustrative example of the FutureSpace Center development investment. For this project, it would be logical to define time T as the end of month 12, the time when the second building is projected to be completed. We would grow the FutureSpace One building's completed value or cash flow forward in time to month 12 at the OCC rate appropriate to the built property (9 percent). We would then combine that with the projected \$5 million value of the FutureSpace Two building as of month 12 to arrive at the total value of V_T , as follows:

$$\begin{aligned} V_T &= \$5,000,000(1.0075)^6 + \$5,000,000 \\ &= \sum_{t=7}^{12} \$37,500(1.0075)^{12-t} + \$10,000,000 = \$10,229,000 \end{aligned}$$

Similarly, with the construction costs, we would grow the four projected payments forward to month 12 at their OCC rate of 3 percent to obtain K_T :

$$\begin{aligned} K_T &= \$1,500,000(1.0025)^9 + \$1,500,000(1.0025)^6 + \$1,500,000(1.0025)^3 + \$1,500,000 \\ &= \$6,068,000 \end{aligned}$$

¹⁰ It can be verified that equation (4) will provide the same OCC for the development (land) expected return as what is implied by the binomial option model presented in Chapter 16 (with period length equal to T) for any state of the world in which it is optimal to develop the property in both of the two subsequent binomial outcomes: in other words, whenever there is a definitive commitment to develop the project. If, contrary to our assumption in this example, it was optimal to wait and not begin immediate development of the FutureSpace site, then equation (4) would imply a value of $E[r_c]$ greater than the true OCC of the land (development option) investment, causing the time 0 valuation implied by equation (4) (the valuation of the formula on either its left-hand side or right-hand side) to be less than the actual land value including the value of the option to wait. This, in turn, would result in a negative NPV calculation for the immediate development project when the opportunity cost of the land is netted out. (See the boxed feature in Section 16.2 entitled: "The NPV Rule and (Re)development Option Value.")

Thus, the projected net development profit as of month 12 is:

$$V_T - K_T = \$10,299,000 - \$6,068,000 = \$4,161,000$$

Substituting in equation (4) and solving for $E[r_C]$ as a simple monthly rate gives:

$$\begin{aligned} \frac{\$10,229,000 - \$6,068,000}{(1 + E[r_C])^{12}} &= \frac{\$10,229,000}{(1.0075)^{12}} + \frac{\$6,068,000}{(1.0025)^{12}} \\ \Rightarrow E[r_C] &= \left[\frac{(\$4,161,000)(1.0075)^{12} (1.0025)^{12}}{(1.0025)^{12} \$10,229,000 - (1.0075)^{12} \$6,068,000} \right]^{\left(\frac{1}{12}\right)} - 1 = 1.54\%/\text{mo} \\ &\Rightarrow (1.01542)^{12} - 1 = 20.16\%/\text{yr} \end{aligned}$$

Thus, the “canonical” OCC of the FutureSpace development investment is 20.16 percent (EAR), which compares to 9.38 percent for the corresponding built property investment in Hereandnow Place and 3.04 percent for risk-free investment. This suggests that the development project has approximately 2.70 [= (20.16% - 3.04%)/(9.38% - 3.04%)] times the investment risk as the built property based on the canonical cash flow timing.

Considering that the difference in risk between the development and built property investments results essentially from leverage, it will not be surprising that the canonical OCC of the development investment can be equivalently derived using the WACC formula we introduced in our discussion of leverage in Chapter 12. If we reproduce the WACC formula of Section 12.3 in its leverage-ratio form and with the variable labels we are using here, the WACC formula for the development OCC is:

Equation 5

$$E[r_C] = E[r_D] + LR(E[r_V] - E[r_D])$$

Defining LR as the leverage ratio, $V/(V - K)$, based on the time 0 valuations of the asset to be built and the land value:

$$LR = \frac{V_0}{V_0 - K_0} = \frac{\$9,352,000}{\$9,352,000 - \$5,889,000} = 2.70$$

we have the development (land) OCC given as:

$$E[r_C] = 3.04\% + 2.70(9.38\% - 3.04\%) = 20.16\%^{11}$$

Note that in whatever manner the canonical OCC is determined, it will, of course, yield the same NPV of the development project (land value) as we derived originally back in equation (2):

¹¹ Recall that the WACC formula is only an approximation for fractional or multiperiod returns; it works exactly only for simple returns, that is, a simple single period ratio, with all cash flows at the beginning and end of the period. While the canonical formula does move all expected cash flows to those two points in time (using appropriate OCC rates), the application of the simple WACC formula to the resulting canonical OCC of the development project will only be exactly consistent with a per annum IRR metric in the special case where the project lasts exactly 1 year (as is true in our FutureSpace illustration).

$$\frac{\$4,161,000}{1.2016} = \frac{\$10,229,000}{1.0938} - \frac{\$6.068,000}{1.0304} = \$9,352,000 - \$5,889,000 = \$3,463,000$$

Prior knowledge of this NPV is not necessary to ascertain the canonical OCC of the development project, as $E[r_c]$ is determined solely by the variables on the right-hand side of equation (4) [or equivalently of (5)].

Exhibit 18-2 graphically portrays what we have just discovered about the relative risk and return for investments in Hereandnow and FutureSpace, using the familiar equilibrium risk/return diagram with the Security Market Line (SML) you have seen so often in this book. Note that the exact outcome percentage ranges on the horizontal axis in the exhibit are just illustrative (see Section 18.2.1 below for their derivation). The point is that the ratio of the amount of risk between the development project and the stabilized asset, as indicated on the horizontal axis, is the 2.70 ratio that we have just computed, the same ratio as between the two expected return risk premia.

18.1.3 THE DEVELOPMENT RISK RATIO AND IMPLICATIONS FOR CHARACTERISTIC DEVELOPMENT REGIONS

The canonical formula for the OCC of a development project, which is most useful in projects that produce a single asset in a single phase, yields an interesting by-product: the **development risk ratio (DRR)**, which is the ratio of the amount of investment risk in the development project to the amount investment risk in an unlevered investment in an otherwise identical stabilized property. As noted above and illustrated in Exhibit 18-2, the DRR in our FutureSpace project example is 2.70, determined as the ratio of the risk premium in the development project canonical OCC (20.16% - 3.04%) divided by the risk premium in the asset being produced by the development project (9.38% -

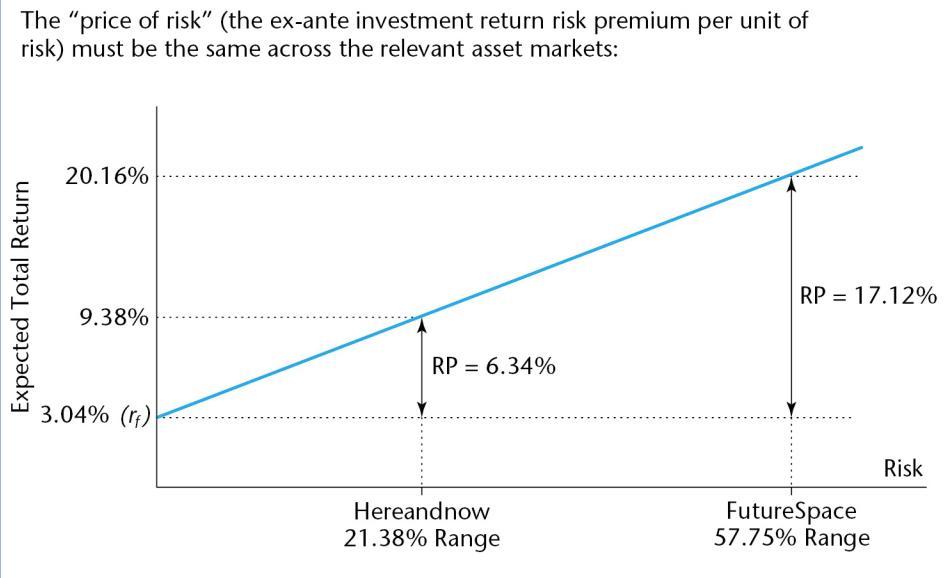


EXHIBIT 18-2 FutureSpace and Hereandnow Risk and Return Based on the Canonical Formula.

If this relationship does not hold, then there are “supernormal” (disequilibrium) profits (expected returns) to be made somewhere and correspondingly “subnormal” profits elsewhere, across the markets for Land, Stabilized Property, and Bonds (“riskless” CFs).

3.04%). As we have seen, this ratio measures the amount of operational leverage in the development project.

We can use the development risk ratio to analyze the characteristic economics of the development project in a little more depth by combining the implications of the canonical OCC analysis with the real options model of land value introduced in Chapter 16 and with the (deterministic) model of optimal density presented in the text box in Section 17.1.1 in the previous chapter. After all, the development project is the (presumably optimal) exercise of such a real option. The project is presumably the HBU of the site, which includes a reflection of the optimal density for the location as determined in the Preliminary Phase of the development process. In the option model, it is greater uncertainty that implies a higher optimal land value fraction and resulting greater density. In the Wheaton-DiPasquale deterministic optimal density model, a higher location value (land value) implies higher optimal density.

As will be shown in the next section, the development project OCC equals the real option (land speculation) OCC at the moment when it is optimal to develop the land (that is, to exercise the real option). Now recall from Chapter 16 that the ratio of the real option OCC risk premium to the underlying asset OCC risk premium is given by the option elasticity (which we labeled η , “eta,” in Section 16.5). Under the assumptions of the Samuelson-McKean formula, this elasticity is a simple function of, among other things, the volatility and the payout ratio (yield rate, similar to the cap rate) of the underlying asset (the property to be built). In particular, greater volatility and/or lower payout (lower cap rate) in the real estate market implies lower option elasticity. Furthermore, the development project’s land value fraction—the fraction of land value in the total value of the to-be-built asset at the time of optimal development—is simply the inverse of the option elasticity.¹² Thus, greater volatility and/or lower cap rates imply a greater land value fraction.

Greater volatility and/or lower cap rates are thus associated with higher land values and less propensity to build. In effect, the developer must hold back and wait for a greater “irreversibility premium,” as this term was defined in Chapter 7. This means that the city with greater volatility and lower cap rates in its real estate market will tend to be denser with higher rents, other things being equal. Higher land value fractions are also associated with lower real depreciation rates for built property assets, as the land value does not generally depreciate, only structure value. This would imply higher nominal appreciation rates in areas with greater volatility and/or lower cap rates. Indeed, lower cap rates would be directly associated with higher asset value appreciation rates, holding constant the risk (and therefore the expected total return) in the built assets.

Greater volatility and/or lower cap rates will also be associated with less operational leverage and, therefore, less additional risk in the development project investment (other things being equal, and as compared to the risk in unlevered investment in the stabilized property). This simply reflects the lower development risk ratio (or, equivalently, the lower option elasticity), which is just the flip side of the coin of the higher land value fraction (since the land value fraction is the inverse of the option elasticity). This means less operational leverage; that is, more of the completed project’s value is in the land’s value relative to the construction cost in high-land-cost locations. However, this does not mean that the absolute physical amount of construction per acre (density or capital intensity) will be less, as high land value generally implies high-density HBUs.¹³ The resulting

¹² Rearranging formula (2c) of Section 16.5.1, we have:

$$V^* = \frac{K_0 \eta}{(\eta - 1)} \Rightarrow \frac{(V^* - K_0)}{V^*} = \frac{1}{\eta}$$

where as in Chapter 16, V^* is the current value of the completed project (at the moment of optimal development), K_0 is the present value of the construction costs, and η is the option elasticity.

¹³ In the optimal density model of the Section 17.1.1 text box, land value or location value is essentially represented by the intercept in the demand function for “P”, the building price per square foot (what we called the “starting sales price”, \$1000/SF in our numerical example depicted in the graph in the text box). It can be verified that, *ceteris paribus*, higher values of that intercept are associated with higher optimal density and higher land value fractions in that model.

implication for construction project duration may counteract to some extent the tendency for high land value to be associated with low development risk ratios, as long construction duration tends to increase operational leverage.

In some cases, the DRR can help flag a particularly risky or aggressive development project, or in the other extreme, one that does not fully use the site's value. For example, suppose you are in a high volatility or low cap rate type of real estate market, that is, the type of market where the option elasticity and therefore the development risk ratio should normally be relatively low (corresponding to a high land value ratio). Yet the development project you are planning seems to pencil out at an (honestly) estimated very high-risk ratio based on the canonical formula. This suggests that the construction project may not add enough value to the site, particularly if it is not a very long-duration, capital-intensive project. The high-risk ratio in the planned project in such circumstances likely reflects that the projected stabilized property value is not sufficiently in excess of the projected construction cost, resulting in present-value profitability, $V_0 - K_0$, which is too low. (Perhaps it is the wrong type or scale of building for the site.) While this might be seen by a direct NPV analysis comparing $V_0 - K_0$ with the land value, it might be difficult to estimate the land value accurately. Alternatively, suppose you are in a low-volatility or high-cap rate real estate market (which normally supports projects with high-risk ratios), but your project presents a low-risk ratio. Then, the structure you are building may be too small or "cheap" for the site. A relatively high $V_0 - K_0$ measure of present-value profitability (which could cause a low-risk ratio) might actually be indicating that your project is relying too much on the (already low) land value and not enough on value added by the structure.¹⁴

*18.2 ADVANCED TOPICS IN DEVELOPMENT

Next, we want to discuss some more advanced topics. First, we will relate the DCF for developments to the real options models from Chapter 16. In Section 18.2.2, we return to the topic of capital stack and its impact on returns (Chapter 15). We already hinted in Chapter 17 that the general/limited partner structure is widespread in development. Finally, in Section 18.2.3, we "unblend" the "development IRR" into the different phases of development, tying it back to Chapter 16. As you probably noticed, these sections are all starred, meaning they are not for everyone.

*18.2.1 THE RELATIONSHIP OF DEVELOPMENT VALUATION TO THE REAL OPTION MODEL OF LAND VALUE

As noted, the development project valuation and relative risk analysis described here are completely consistent with the real options model of land value described in Chapter 16. We are here assuming that the point has arrived where it is optimal to make an irreversible commitment to construct the entire development project. In this context, the canonical formula demonstrates the equivalence of the approach described here and the real options valuation model of Chapter 16.¹⁵

¹⁴ Again, an accurate assessment of land value would presumably cause the NPV of a low risk ratio project in a low land value region (where risk ratios should be higher) to appear negative (net of land cost), as the small scale or cheap nature of the project (causing the K_0 value to be small) would cause the $V_0 - K_0$ value to be lower than the land value. But the land value may be difficult to accurately estimate independent of the proposed project's economics (the $V_0 - K_0$ value itself). So the risk ratio may be a telling indicator. However, please note that this type of use of the DRR metric is only suggestive. It may flag situations that deserve further looking into, but does not necessarily imply a suboptimal development plan or timing. Ultimately the NPV Decision Rule (applied net of true land opportunity cost) should be the guiding principle.

¹⁵ Note that the NPV and real option models presented here assume that the development project involves a single phase, effectively a single asset (even if it might be several buildings) that are committed irreversibly to be entirely completed once the decision to commence construction is made. More complex development projects that involve two or more sequential phases, in which later phases retain flexibility in timing, magnitude and design, or even whether to be built at

To see this, consider again our FutureSpace Center example. Suppose that risk in the market for stabilized office property can be thought of as making the expected development asset value upon completion, what we have quantified as $V_T = \$10,229,000$, a random variable that has the following possible outcomes at time T :

$$V_T = \$11,229,000; \text{ or}$$

$$V_T = \$9,229,000.$$

With an ex-ante probability of 50 percent for each outcome. Then, applying the certainty-equivalence form of the DCF present value model that we first introduced in Chapter 16 and which formed the basis of the binomial option value model, we have the following present value computation for the FutureSpace project as of time 0:¹⁶

$$PV[C_1] = \frac{E_0[C_1] - (C_1^{up} \$ - C_1^{down} \$) \frac{E_0[r_v - r_f]}{V_1^{up \%} - V_1^{down \%}}}{1 + r_f}$$

$$\begin{aligned} PV[C_1] &= \\ & \frac{((0.5)(\$11.229 - \$6.068) + (0.5)(\$9.229 - \$6.068)) - \left(\frac{(\$11.229 - \$6.068) - (\$9.229 - \$6.068)}{\$11.229 - \$9.229} \right) \frac{9.38\% - 3.04\%}{\$9.352}}{1.0304} \\ &= \\ & \frac{\$4.161 - \$2.000 \left(\frac{6.35\%}{120.07\% - 98.69\%} \right)}{1.0304} = \frac{\$4.161 - \$2.000(0.2964)}{1.0304} = \frac{\$3.568}{1.0304} = \$3.463 \end{aligned}$$

As the development project is irreversibly committed, the binomial model reflects option exercise in either of the two future outcomes. Thus, the certainty-equivalence binomial option models yield the same \$3.463 million valuation of the project and, therefore, the same implied OCC of 20.16 percent, as we derived previously:

$$1 + E[r_c] = \frac{E[C_T]}{PV[C_T]} = \frac{\$4.161}{\$3.463} = 1.2016$$

Note that in percentage terms, the investment outcome spread (which may be viewed as a measure of risk) is:

all, require in principle more complex analysis and evaluation models. For a more practical approach to analyzing such projects using Monte Carlo simulation, see: Geltner, D. & R. de Neufville, *Flexibility and Real Estate Valuation Under Uncertainty: A Practical Guide for Developers*, Wiley Blackwell, Oxford, 2018.

¹⁶ Recall that: $V_T^{up \%} - V_T^{down \%} = (V_T^{up \$} - V_T^{down \$}) / PV[V_T]$ and that $PV[V_T]$ is $\$10,229,000 / 1.0938 = \$9,352,000$.

$$\frac{\$5,161,000 - \$3,161,000}{\$3,643,000} = \frac{\$2,000,000}{\$3,643,000} = 57.75\%$$

in the FutureSpace development project, whereas for this project's "underlying asset" (the assets to be built), the percentage investment outcome spread is only:

$$\frac{\$11,229,000 - \$9,229,000}{\$9,352,000} = \frac{\$2,000,000}{\$9,352,000} = 21.39\%$$

The ratio of these outcome spreads (risk): $57.75/21.39 = 2.70$, is exactly the same as the ratio of the ex-ante risk premia between the development project and its underlying asset: $(20.16\% - 3.04\%)/(9.38\% - 3.04\%) = 17.12\%/6.34\% = 2.70$. This demonstrates our previous point that the project valuation model presented here is based on the assumption that the "price of risk" (the ex-ante investment return risk premium per unit of risk) must be the same across the relevant asset markets, in this case, the market for development projects (land) and the market for the stabilized assets those projects produce.

*18.2.2 CAPITAL STRUCTURE FOR REAL ESTATE DEVELOPMENT

In Chapter 15, we introduced the concept of a "capital stack." More specifically, we showed how an investment can have multiple investors organized in a joint venture. In short, we found that the general partner (GP) does most of the preliminary phase and operational management of the project. In contrast, the limited partner (LP) pays most of the bills while retaining primary control over major capital decisions. The GP tends to be a local, entrepreneurial investor, whereas the LP tends to be larger institutional investors or wealthy individuals. The incentive structure works because the GP can achieve high returns if the investment performs really well. This will benefit the LP as well.

In this section, we want to simplify the example in Chapter 15 for development. This is important, as most developments will be organized in some sort of joint venture structure. The reason is that developments are generally very local projects. You need to be "in the know" to be successful in this business. Many large corporations do not have the time or appetite for the risk associated with managing the preliminary stage of development; see Chapter 17 as well. Typically, a GP will find a piece of land with (hopefully) much potential. The GP will make the designs, get the permits, etc. Once this phase is done, and it is time to purchase the land or commit to construction, the GP will "shop" for an LP willing to finance the equity component of the project. The LP may largely finance the land purchase if the developer does not already own the land and help to arrange the construction loan (often, the GP does not have enough assets/history to procure such a loan). During the construction phase (phase 2 in Exhibit 17-1), the GP will act as the project manager to ensure the development is on time and on budget and undertake the marketing and leasing of the project. The GP will also get a fee for his role as project manager (and, to some extent, to cover overhead or past preliminary phase expenditures). Once the development is done, the investors can either sell the property or retain it, and the two partners often retain flexibility for either or both to sell out at that time by means of a "buy/sell agreement."¹⁷

Exhibit 18-3 shows the proforma for a property worth \$10 million that would generate \$400,000 in cash flow (property level) at the end of year 0. Both the cash flows and the property value grow by 3% annually. (See Panel A.) This results in a going-in cash flow yield of 4.12% ($\$412,000/\$10,000,000 = 4.12\%$).

¹⁷ In essence, in a typical buy/sell agreement, either party can buy out the other party by offering a price. If the other party does not want to sell at that price, then they must buy from the first (offering) party at that same price.

EXHIBIT 18-3

Example of Capital Structure for a Real Estate Development

A. Similar Stabilized Rental Property							
Year	0	1	2	3	4	5	6
1 Property value	\$10,000,000	\$10,300,000	\$10,609,000	\$10,927,270	\$11,255,088	\$11,592,741	\$11,940,523
2 PBTDCF	\$ 400,000	\$ 412,000	\$ 424,360	\$ 437,091	\$ 450,204	\$ 463,710	\$ 477,621
B. Permanent Mortgage							
Year	0	1	2	3	4	5	6
3 Loan amount	\$ -	\$ -	\$ 8,000,000	\$ 8,000,000	\$ 8,000,000	\$ 8,000,000	\$ 8,000,000
4 Debt service	\$ -	\$ -	\$ 400,000	\$ 400,000	\$ 400,000	\$ 400,000	\$ 400,000
5 EBTCF	\$ -	\$ -	\$ 24,360	\$ 37,091	\$ 50,204	\$ 63,710	\$ 77,621
C. Development							
Year	0	1	2	3	4	5	6
6 Land acquisition cost	\$ 2,000,000	\$ -	\$ -	\$ -	\$ -	\$ -	\$ -
7 Construction cost	\$ -	\$ 8,000,000	\$ -	\$ -	\$ -	\$ -	\$ -
8 Construction loan	\$ -	\$ 8,000,000	\$ -	\$ -	\$ -	\$ -	\$ -
9 Construction loan repayment	\$ -	\$ (8,000,000)	\$ -	\$ -	\$ -	\$ -	\$ -
D. Preferred Equity							
Year	0	1	2	3	4	5	6
10 Equity LP (begin)	\$ -	\$ 1,800,000	\$ 1,899,000	\$ 1,979,085	\$ 2,050,844	\$ 2,113,437	\$ 2,165,966
11 LP Investment	\$ 1,800,000	\$ -	\$ -	\$ -	\$ -	\$ -	\$ -
12 Preferred equity (LP)	\$ -	\$ 99,000	\$ 104,445	\$ 108,850	\$ 112,796	\$ 116,239	\$ 119,128
13 Preferred equity (paid)	\$ -	\$ -	\$ 24,360	\$ 37,091	\$ 50,204	\$ 63,710	\$ 77,621
14 Preferred equity (Unpaid)	\$ -	\$ 99,000	\$ 80,085	\$ 71,759	\$ 62,593	\$ 52,529	\$ 41,507
15 Pay off accrued LP balance	\$ -	\$ -	\$ -	\$ -	\$ -	\$ -	\$ -
16 Equity LP (end)	\$ 1,800,000	\$ 1,899,000	\$ 1,979,085	\$ 2,050,844	\$ 2,113,437	\$ 2,165,966	\$ 2,207,473

\$10,000,000) of this reference stabilized property. Obviously, this cash cannot be collected from the subject development property, as that property has yet to be developed. The land is purchased for \$2 million at the end of year 0, and the construction is finished at the end of year 1. A construction loan of \$8 million is procured to finance the construction. All construction costs are paid for by this loan (including the interest). (Panel C.) Once the property is developed, an (interest-only) permanent loan replaces the construction loan. The interest on this loan is assumed to be 4%. (See Panel B.) The LP takes on 90% of the costs and, in return, will enjoy a 5.5% preferred return. Once the LP gets its preferred return, any remaining spoils are split 50/50 between the LP and GP. The same goes for the proceeds when selling the property after holding it for 5 years in the current example. This structure is similar to that in Chapter 15.

Thus, in year 0, the LP will invest \$1.8 million, and the GP will invest the remaining \$200,000. In year 0, there is still no income as the property is still in development. As a result, the LP gets none of its preferred returns, which will accrue into its capital account. We can track the capital account in Panel D of Exhibit 18-3. Starting in year 2, the property will generate cash flows. However, the EBTCF is never enough to hit the preferred return fully. More details on capital accounts are provided in Chapter 15. You can also play around with the Excel file online to track our steps in more detail.

Exhibit 18-4 provides the property and partners' internal rate of return. Furthermore, we also provide the IRR separately for the construction phase (i.e., you sell immediately after the development is finished) and the construction plus operational phase. If we start with the PBTDCF, we

EXHIBIT 18-4

IRR on Investment for Partners and Different Phases

Internet Rate of Return

Phase	IRR	0	1	2	3	4	5	6
PBCF Both	7.46%	\$ (2,000,000)	\$ (8,000,000)	\$ 424,360	\$ 437,091	\$ 450,204	\$ 463,710	\$ 12,418,144
	Construction	15.00%	\$ (2,000,000)	\$ 2,300,000				
	Operational	7.12%		\$ (10,300,000)	\$ 424,360	\$ 437,091	\$ 450,204	\$ 463,710
EBTCF Both	13.38%	\$ (2,000,000)	\$ -	\$ 24,360	\$ 37,091	\$ 50,204	\$ 63,710	\$ 4,018,144
	Construction	15.00%	\$ (2,000,000)	\$ 2,300,000				
	Operational	13.05%		\$ (2,300,000)	\$ 24,360	\$ 37,091	\$ 50,204	\$ 63,710
LP	Both	11.60%	\$ (1,800,000)	\$ -	\$ 24,360	\$ 37,091	\$ 50,204	\$ 63,710
	Construction	11.08%	\$ (1,800,000)	\$ 1,999,500				
GP	Both	25.00%	\$ (200,000)	\$ -	\$ -	\$ -	\$ -	\$ 762,788
	Construction	50.25%	\$ (200,000)	\$ 300,500				

invest \$2 million in year 0 and receive \$2.3 million in year 1 after selling the property. The IRR of the construction phase is therefore determined to be 15%. This is about twice the IRR of the investment if held for the full 6 years. This reveals how the construction phase is more risky compared to the operational phase due to the effect of operational leverage (at this PBCF level, there is no financial leverage, and there is no lease-up risk in this project since the cash flows match those of the Panel A stabilized twin asset). If we turn our attention to the LP during the construction phase, we find he invests \$1.8 million in year 0 while receiving almost \$1,999,500 in year 1. This is computed as the LP's equity balance at the end of year 1 (line 16 in Exhibit 18-3) plus 50% of the residual. The equity balance is the invested amount plus its preferred return in that year: $\$1,800,000 + \$99,000 = \$1,899,000$. The residual, in turn, is computed by subtracting both the LP's equity balance (end of year 1) and the GP's investment from the net proceeds of selling the property immediately after development (after paying back the \$8 million loan). Thus, in addition to the \$1,899,000, the LP receives 50% of $\$10,300,000 - \$8,000,000 - \$1,899,000 - \$200,000 = \$201,000$, or \$100,500. The LP's IRR during the construction phase is now 11.08%, which is very similar to the LP's IRR when holding the property over the entire 6 years (11.6%) and lower compared to the asset-level IRR for the construction phase (15%). The GP in this example joint venture structure presumably faces more risk, as suggested by the greater expected return he faces: 50.25% for the construction phase (see Exhibit 18-4). But we cannot ascertain from these numbers alone if the structure is a "fair" one. For this, you would have to do a scenario analysis similar to that in Chapter 15.

*18.2.3 "SOUP TO NUTS": INCLUDING LEASE-UP AND LAND ASSEMBLY: "UNBLENDING" THE BLENDED IRR

Although our primary focus in this chapter is on the major capital investment phase of development, involving the irreversible commitment of the land to the construction phase, it is of interest now to step back to the question of how the framework and analysis presented in this chapter relate to the overall development process described in the first section of Chapter 17. There, we suggested that there are several stages in most development projects, with different risk and return regimes across the stages. The basic NPV approach presented here can be applied to any one of the phases of the overall development process, not just to the construction phase considered here, and in fact, it should be applied separately to each of the phases. The reason is that each phase is characterized by different levels and types of investment risk, and is characteristically funded by different sources of capital. Nothing requires the investors in one phase to remain in for subsequent phases. Therefore,

it makes more sense to evaluate each phase separately, with OCC discount rates appropriate to each phase. It can muddy the waters to blend the NPV analysis across phases using a single blended discount rate. Nevertheless, that is the typical current industry practice. In this section we describe a way to “unblend” those blended OCC rates typically currently used.

We have repeatedly noted that the project evaluation implied by the methodology developed in Section 18.1.1 is based on the concepts of wealth maximization and equilibrium within and across the relevant asset markets. Assuming that successful developers do indeed maximize their wealth and that real estate markets function reasonably well, then the prices of development projects (or equivalently, the price of land ripe for development) that we observe in the world should tend to conform to the normative values implied by the methodology of Section 18.1.1. To the extent that this is true, no matter which rules of thumb developers apply in the real world (Section 17.4), the successful bidders must end up applying those rules in such a way that they end up with the values and investment decision results implied by the Section 18.1.1 methodology. At least this must hold typically and on average.

For example, in the case of the FutureSpace Project, which has, as we have seen, a normative value of \$3.463 million, the developer would have to apply a profit margin criterion of 5.7%¹⁸ or an enhanced cap rate on cost of 9.69 percent¹⁹ (that is, 69 basis points of enhancement instead of 100), or a blended long-run IRR of 9.72 percent (that is, 72 basis points of premium over the stabilized 9 percent rate) to get to the same \$3.463 million valuation. Therefore, assuming that markets are functioning reasonably well, we can use typical developer rules of thumb for typical project cash flow profiles to “back out” the market’s true OCC rates for each phase of the overall development process that we elaborated in Section 17.1 in Chapter 17.

To illustrate this procedure, we can consider an application to midsized developers of mixed-income rental housing in the Boston metropolitan area.²⁰ Cash flows for all development and stabilized operation phases were estimated for typical mixed-income rental housing projects in the Boston area from the preliminary phase through the stabilized operating phase. Developers were then asked what long-run going-IRR premium over stabilized property IRRs they would require if they entered the project at either the permitting, construction, or lease-up phase. In other words, developers were asked to quantify their decision-making based on a common rule of thumb they were comfortable with, namely what we previously termed the blended long-run IRR approach.

By working backward (starting out with the implied stabilized asset value), computing the value of the development asset at the beginning and end of each phase of the development process based on the cash flows within each and all subsequent phases and on the long-run IRRs stated by the developers, the IRR implied specifically *within* each phase was determined.²¹ The result is presented in Exhibit 18-5. The study, which was done in 2005 when Treasuries were yielding around 4 percent, found that for developments producing properties that had stabilized investment required return expectations of 8.3 percent (going-in IRR based on a 10-year hold), the lease-up phase investment

¹⁸ Computed as $\frac{\$10,000,000}{1.057} - \$6,000,000 = \$3.46 \text{ million}$ (rounded).

¹⁹ Computed as $\frac{\$900,000}{0.0951} - \$6,000,000 = \$3.46 \text{ million}$ (rounded).

²⁰ See W.T. McGrath, “Financial Economics and 40B: A Framework to Mediate the Fight over Money,” CRE working paper, 2005.

²¹ For example, if this procedure were applied to our illustrative FutureSpace Project, developers would presumably state a required long-run IRR rate of 9% for the stabilized investment phase and 9.72% for the development investment phase (noting that this preleased project has no lease-up phase). The stated required 9% stabilized investment rate would then be applied to determine the completed project value of \$5 million for each building, and the stated required 9.72% development project rate would be applied to determine the \$3.463 million value of the construction phase investment. These two values in combination with the construction phase cash flows would then allow us to “back out” the implied development phase IRR that is required by the developers, which we saw to be 20.16%, using the canonical cash flow assumption.

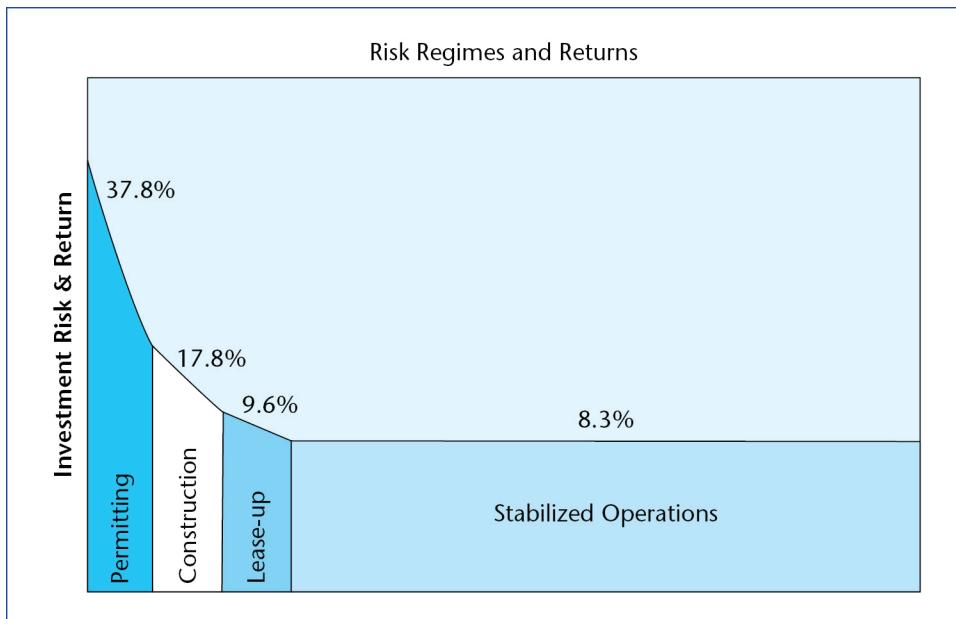


EXHIBIT 18-5 An Example of Risk and Return Regimes for Phases of the Development Process (taken from a study of affordable housing development in Massachusetts).

alone would require approximately a 9.6 percent going-in IRR, the construction phase investment alone would require a 17.8 percent going-in IRR,²² and the preliminary phase investment (land assembly and permitting) would require a 37.8 percent going-in IRR.²³

18.3 CHAPTER SUMMARY

This chapter is built on the theoretical treatment of land value and development timing presented in Chapter 16 to present an overview of the development decision-making process and a rigorous yet practical methodology for the investment analysis and evaluation of development projects. The main point of this chapter was to present and advocate the use of an NPV-based financial evaluation procedure for real estate development projects. The recommended procedure is based soundly on fundamental economic principles of opportunity cost, market equilibrium, and wealth maximization. It is consistent with modern corporate capital budgeting and securities investment theory and the valuation methodology presented in Parts I and III of this book. It requires no more, or no more difficult to obtain, information than the ad hoc rules of thumb used widely in the development industry.

²² Based on the “canonical” cash flow timing assumption and a 1-year construction phase.

²³ As we noted in Section 18.1, relatively little financial capital is typically invested during the preliminary phase (labeled “Permitting” in Exhibit 18-5). Much of the “investment” in this phase is the application of the human resource of the developer’s “entrepreneurial labor,” which may appear as part of the routine operating expenses of the development firm (some or all of which may be recouped in a “Developer’s Fee”). To the extent that there is an investment of financial capital, much of this might be in the form of land speculation (discussed in general from an investment perspective in Chapter 16). However, the developer may add additional value in the preliminary phase, as assembly, permitting, and project design adds value to the site. In the case of the Boston mixed-income housing development study, the investment outlays in the preliminary phase consisted primarily of nonrefundable land optioning payments plus direct expenses of project permitting and design. These preliminary phase investments totaled only about 3% of the subsequent phase’s land acquisition and construction cost investments.

The chapter also discussed some ramifications, implications, and information byproducts of the canonical development project NPV procedure, which have interesting practical and academic uses.

KEY TERMS

- Financial desirability
- Time-to-build
- Phased risk regimes
- Operational leverage
- Price of risk
- Canonical formula of development project OCC
- Real options model of land value
- Developmental risk ratio
- Profit margin ratio
- Enhanced cap rate on cost
- Blended long-run IRR
- IRR “unblending”

STUDY QUESTIONS

Conceptual Questions

- 18.1. Describe the three features that characterize real estate development projects and distinguish them from investments in fully operational (aka “stabilized”) properties.
- 18.2. What is *operational leverage* in the context of a real estate development project? Explain why this characteristic causes investment in a development project to be riskier than investment in an otherwise similar stabilized property.
- 18.3. Explain why the opportunity cost of capital (OCC) of construction costs is usually not much greater than the risk-free interest rate.
- 18.4. Explain why a lender’s expected return on a construction loan is lower than the stated construction loan interest rate and generally not much higher than the risk-free rate.
- 18.5. What two points in time are useful to define for the purpose of applying the NPV investment rule to development projects and for quantifying the development phase OCC?
- 18.6. Explain the economic intuition underlying the “canonical” formula for the development phase opportunity cost of capital (OCC), $E[r_c]$, in equation (4).
- 18.7. Under what condition is equation (4) consistent with the real option model of land value and optimal development timing of Section 16.3?
- 18.8. In general, what is the relationship between the highest and best use of a vacant site, the NPV of development exclusive of land cost, and the land value, assuming the site is ripe for development?
- 18.9. Describe the “enhanced cap rate on cost” technique developers often employ to examine the financial viability of development projects. How can this approach be consistent with the economically rigorous evaluation methodology developed in this chapter?

Quantitative Problems

- 18.10. (NPV-Based Front Door & Back Door) This problem revisits the proposed 20,000 SF office/warehouse flex space development first introduced in Problem 17.9 and gets you to apply an NPV-based back-door financial evaluation related to the approach recommended in this chapter. Market rents are about \$15/SF and the local area has a 5% vacancy rate.

All operating expenses are passed through to tenants except property taxes, insurance, and management, which you estimate at \$5/SF per year. Construction costs are estimated at \$1,030,075 as of project completion in 1 year (assume that construction is riskless). Lease-up is assumed to be instantaneous and the projected cap rate on the completed stabilized property NOI is 10%. The opportunity cost of capital for stabilized property is 11% and the risk-free rate is 6%.

- a. Use an NPV-based back-door approach to estimate the maximum land acquisition costs that can be supported by this project, assuming there are no other upfront development costs or fees other than land acquisition. (Hint: Determine the time 0 value of the developer's NPV exclusive of land cost as in the right-hand side of equation (4).)
- b. How does your answer in part (a) compare with the answer to Problem 17.9 part (a)? Which one is the better indicator of the true economic value of the site, and why? What information is used in the Chapter 17 SFFA approach that is not used in the NPV-based version, and vice versa?
- c. Now apply an NPV-based front-door investment evaluation of the development and derive the required rent assuming the upfront land acquisition cost (inclusive of any other up-front development fees) is \$700,000, and that this figure represents the developer's time 0 equity investment. (Steps: Grow the time 0 investment using the development phase OCC from the left-hand side of equation (4) to determine the minimum required value of the developer's equity as of the time of completion. Combine this with the construction cost as of completion to obtain the required projected stabilized asset value as of time of completion. Apply the cap rate to the result to obtain the minimum required NOI, and then proceed as you did with the SFFA version.)

18.11. A developer is evaluating the economic profitability of a large-scale shopping center that will take an estimated two years to complete. Total (nonland) construction costs compounded forward to the time of completion are estimated to be \$95 million and the estimated value of the completed stabilized property is \$128 million. If the OCC of the construction costs is 5% and the OCC (going-in unlevered IRR) of the completed stabilized property is 8.5%, then, according to the equilibrium relationship in equation (4), what is the appropriate opportunity cost of the development phase investment? (Assume annual compounding.)

18.12. Consider the development of a 100% preleased office building. Total development costs are estimated to be \$55 million. The current (time 0) opportunity cost, or market value, of the land is \$15 million and construction cost draws are projected to be \$40 million in total occurring as follows: \$10 million at time 0, \$20 million 1 year later at the end of year 1, and a final \$10 million cost at the time of project completion at the end of year 2. The value of the completed stabilized property at the end of year 2 (VT) is projected to be \$65 million. T-bills are currently yielding 5.5%. The opportunity cost of capital (OCC) for construction costs contains a 50-basis-point risk premium over T-bills and is at 6%. Stabilized built property commands an investment risk premium of 300 basis points over T-bills in its going-in IRR. With this information in mind, answer the following questions related to this development project. (Note, it is recommended that you answer this in an Excel worksheet so you can go back and examine the impact of changes in key variables (e.g., OCCs and property value at completion).)

- a. Use the OCC of the construction costs to calculate the present value of the expected construction costs as of time 0, K_0 , and use the OCC of stabilized built property to determine the present value of the \$65 million estimated value upon completion, V_0 .
- b. Use your calculations from part (a) to compute the time 0 NPV of the development project, ignoring any upfront (time 0) costs other than the land [and the initial

construction cost draw, of course, which is already accounted for in your calculation in part (a)].

- c. Suppose there is \$1 million of additional upfront fees and costs, besides the land. Based on your answer in part (b), what is the implied maximum price the developer could pay for the land such that this development project would still be desirable at the present time?
- d. Use your answer from part (c) along with the development cash flows (construction costs & stabilized value) to calculate the OCC of the development phase cash flows. (Hint: Find the IRR or discount rate that sets the NPV of development cash flows, including, upfront land and other costs, to zero.)
- e. Assuming that \$15 million is the true current market value of the land (and that other up-front development fees would be no more than the \$1 million), and assuming that the \$40 million cost and \$65 million value projections are realistic, what does your analysis in part (c) tell you about the relative ability of this developer to use this land parcel as compared to other developers?
- f. If the subject developer is the first-best developer of this site, then is it possible that immediate initiation of the \$40 million construction project might not be optimal for this developer? Explain.

18.13. (NPV of Development with a Lease-Up Phase) Consider a speculative development project with the following characteristics: The current (time 0) market value of the land is \$2 million and there are \$200,000 in up-front design fees and developer costs attributable to time 0. Construction is expected to take three years, and the construction contractor is to be paid in three equal annual payments estimated to be \$1.5 million at the end of year 1 through 3. While the exact amounts may vary ex-post due to change orders, these costs are largely fixed by the pre-existing physical design of the structure. A construction loan will be obtained at a projected interest rate of 7.5% plus a \$20,000 up-front origination fee (assume the construction loan is riskless so that $E[r_D] = r_D$). The construction loan will cover all of the projected \$4.5 million direct construction cost, and the loan will be due in its entirety at the end of year 3. The opportunity cost of equity capital invested in the development phase is estimated to be 20% per annum. The construction phase is expected to be followed by two years of absorption, with two annual net cash flows during this phase, at the end of years 4 and 5. The expected net cash flow during year 4 is -\$100,000, as the building is expected still to be largely empty and to incur substantial tenant build-out and leasing expenses during that year. The expected cash flow in year 5 is+\$400,000, reflecting more rental revenue and fewer leasing expenses as the building fills up that year. This lease-up phase is less risky than the development phase, but more risky than the stabilized operating phase. The OCC for the absorption phase is estimated at 300 basis points higher than that for the fully operational property (details on this below). The lease-up phase is followed by the long-term fully operational phase of the completed project, starting from year 6 and extending indefinitely into the future. Net property cash flows from this phase are expected to start out at \$800,000 per year and grow annually each year thereafter at a rate of 1% per year. From the perspective of the beginning of year 6, when the building has first become stabilized, these expected future cash flows have the risk of an unlevered investment in a fully-operational property. This is an amount of risk that warrants a 9% discount rate (going-in IRR from that time forward). *Determine the time zero NPV of the proposed development.* (Hint: Work backwards in time, starting with the stabilized property value after lease-up is completed, V_5 . Use this value along with the expected cash absorption/lease-up year cash flows to calculate the estimated property value at the end of the construction period, V_3 . Be sure to use the appropriate phase-specific OCCs.)



Taylor & Francis
Taylor & Francis Group
<http://taylorandfrancis.com>

Part V

Mortgages from an Investment Perspective

In Part III, we discussed real estate debt from the perspective of the equity investor, the borrower. We pointed out that, from the broader perspective of the capital markets and the investment industry as a whole, the borrowing transaction creates two different types of “investment products” from the same underlying real asset: levered property equity and debt secured by the property. Each of these two products appeals to different types of investors. Investors who are more risk-averse, possibly tolerant of inflation risk, and who need more current income from their investments, or who want an asset with a fixed schedule of future payments and do not want to be involved in property management, find real estate debt to be an appealing type of investment product. The debt market equates the supply (borrowers wanting debt capital) and demand (debt investors) in the market for debt products of all types, including mortgages, which are debt products backed by real estate assets. (Or one can look at it the other way, the borrowers representing the “demand” for financial capital—money—while the lenders provide the “supply” of such capital.)

The debt market as a whole is arguably the largest, most complex, and most technically sophisticated of all the branches of the capital markets. It includes major components such as short- and long-term U.S. Treasury securities, international bonds, corporate and municipal bonds, and, importantly, securities backed by both residential and commercial mortgages. In the United States alone, by the early 2020s, there were over \$53.3 trillion worth of debt securities outstanding, approximately 30 percent of which were based directly or indirectly on mortgages. During the COVID-19 pandemic, mortgage-backed securities (MBS) surpassed U.S. Treasury securities as the largest fixed-income issuance, but their volume quickly declined with the Fed’s rate hikes beginning in early 2022.

Even if your primary focus is on real estate equity investments, it’s essential to have a solid understanding of the debt side, particularly from the perspective of the real estate debt investor. This perspective not only deepens and rounds out the equity investor’s knowledge but also opens up many professional career opportunities. For many years, Wall Street investors who simultaneously invest in both real estate equity and debt have become key players in the real estate market.

With this in mind, in Part V, the following three chapters will introduce the core concepts, terminology, and methodologies essential for understanding real estate debt from the perspective of a debt investor. This section has two primary objectives. First, it aims to equip you with a solid understanding of the debt investor's perspective, enabling you, as an equity investor or in collaboration with one, to effectively navigate deals with debt investors. Second, it seeks to broaden your knowledge of real estate debt, empowering you to creatively structure financing packages (aka "the capital stack"), whether you are working from the equity or debt side. Chapters 19 and 20 will present mortgage fundamentals (most of which are in fact relevant to both residential and commercial loans). Chapter 21 will then delve in more depth into the underlying economics of the debt market in general, and commercial mortgage investment and underwriting in particular.

19 Mortgage Basics I

An Introduction and Overview

CHAPTER OUTLINE

- 19.1 Schematic Typology of Mortgages and Overview of the U.S. Mortgage Industry
- 19.2 Basic Legal Concepts and Terminology for Mortgages
 - 19.2.1 Legal Structure of Mortgages
 - 19.2.2 Priority of Claims in Foreclosure
 - 19.2.3 Typical Mortgage Covenants
- 19.3 Default, Foreclosure, and Bankruptcy
- 19.4 Nonlitigious Actions and Workouts
 - 19.4.1 Costs of Foreclosure
 - 19.4.2 Nonlitigious Actions
 - 19.4.3 Workouts
- 19.5 Limited Liability and Foreclosure Costs: The Borrower's Put Option and Strategic Default
 - 19.5.1 Limited Liability and the Borrower's Put Option
 - 19.5.2 Foreclosure Costs and Strategic Default
- 19.6 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The legal and financial structure of mortgages.
- The major different types of real estate mortgages in the United States and how the mortgage industry works.
- The major legal terms and legal characteristics of mortgages.
- The nature and costs of the foreclosure process and the workout process as a way to avoid foreclosure.

The classical form of real estate debt is the **mortgage**, a loan secured by real property as collateral. The word *mortgage* comes from two Middle English words (which are actually French in origin): *gage* meant an obligation or commitment (as in our modern word *engaged*), while *mort* referred to death or dying. Hence, it was a “dying commitment,” that is, a commitment that was not permanent but had a finite lifetime. In return for present value obtained, a landowner committed himself to pay compensation in money or in-kind over a period of time in the future, but not forever. As land was the direct source of most wealth in medieval times, it was natural to secure the mortgage with a pledge of real property. Thus, mortgage lending is one of the oldest forms of debt recognized in Anglo-Saxon law.

In this chapter, we will focus on the basic legal and financial structure of mortgages, covering fundamental terminology and concepts that apply to all types of mortgages, whether residential or commercial, and to both standard fixed-rate and adjustable-rate loans, as well as other more specialized types. While our primary focus is on loans for commercial property, the foundational

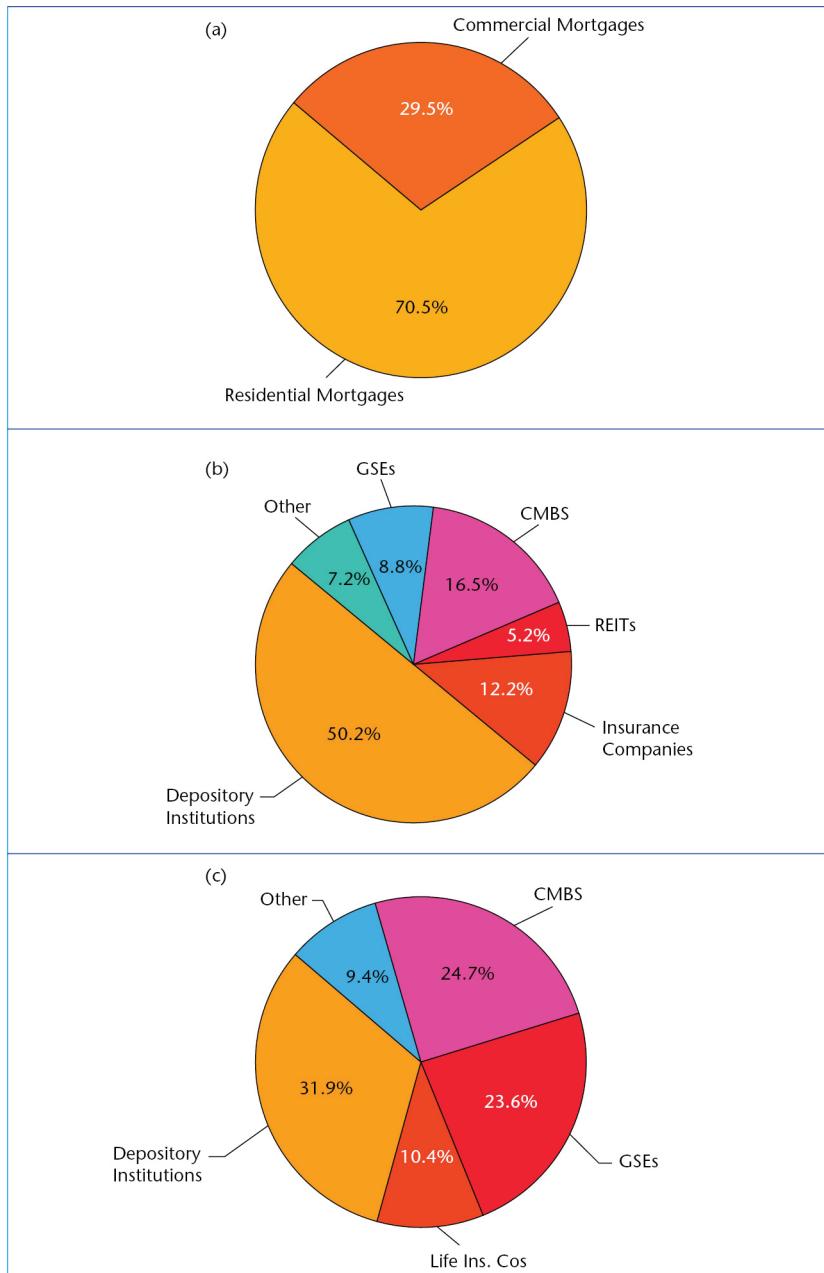


EXHIBIT 19-1 Overview of U.S. Mortgage Debt Outstanding as of 2023.

Panel A U.S. Mortgages Outstanding, 2023 \$19.8 trillion.

Source: FRB.

Panel B U.S. Commercial Mortgages Sources, 2023 \$5.9 trillion outstanding.

Source: FRB.

Panel C U.S. Multifamily Mortgages Sources, 2023 \$2.1 trillion outstanding.

Source: FRB.

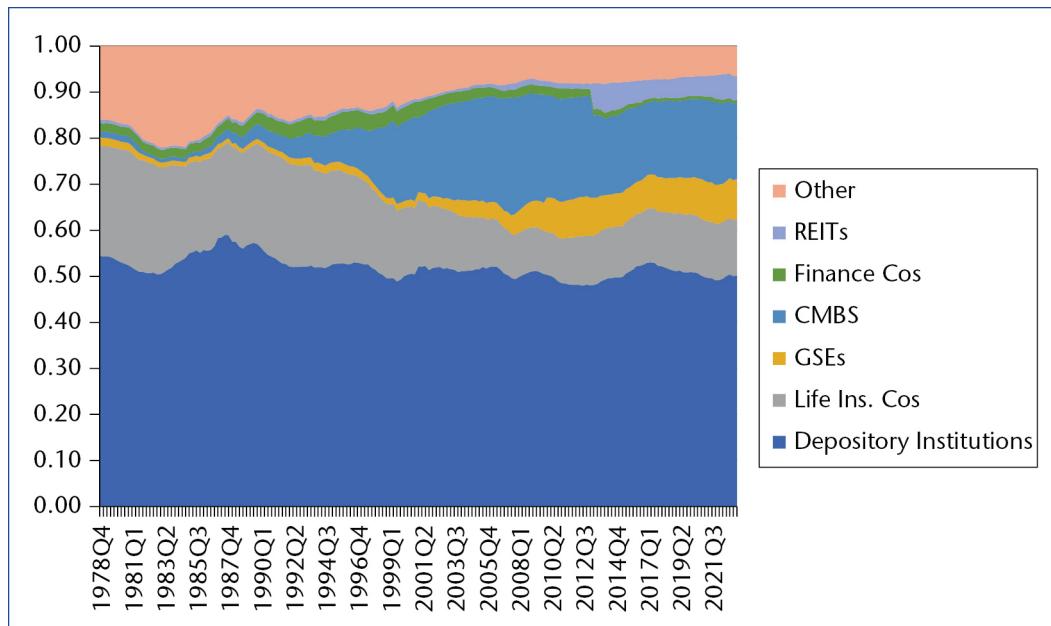


EXHIBIT 19-2 U.S. Commercial and Multifamily Mortgages Share of Outstanding Balance Held by Various Capital Sources, 1978–2023.

Source: Federal Reserve Board.

aspects discussed here are common to all mortgages. We'll start with an overview of the different types of mortgages and branches of the mortgage industry.

To begin, it is important to recognize the magnitude and significance of mortgage lending in real estate investment. Mortgages have traditionally been the largest source of financial capital for most real estate purchases and investments in the United States. Exhibit 19-1 provides an overview as of 2023. Panel A shows that more than 70 percent of the \$19.8 trillion U.S. mortgage debt outstanding was secured by owner-occupied homes, with \$5.9 trillion of commercial mortgages outstanding.

Panel B of Exhibit 19-1 shows the sources of the capital as measured by who is holding the \$5.9 trillion of commercial mortgages outstanding as of 2023. Exhibit 19-2 shows the evolution of the capital sources since 1979. The largest source includes depository institutions, including many small to medium sized local and regional banks. This concentration of capital among depository institutions is closely linked to the **Community Reinvestment Act (CRA)**, which was enacted by Congress in 1977 to encourage depository institutions to help meet the credit needs of the communities in which they operate. The CRA aims to ensure that banks serve the needs of the entire community, including smaller, underserved markets. Local depository institutions typically specialize in construction loans and short-term loans on completed properties. Construction and short-term commercial mortgages are often part of the banks' traditional "relationship lending" business, in which the bank cultivates a local business as a client for various services including as a borrower from the bank, and the bank then tends to hold such loans in the bank's own asset portfolio. Banks and such depository institutions tend to be particularly important in lending to smaller commercial properties and those in smaller communities. Nonbank sources of capital are more important for the larger properties.

Life insurance companies (LICs) are also a major source of commercial mortgages, particularly for larger properties and longer-term loans, typically ranging from 5 to 10 years in maturity. The LIC-sourced component in Exhibit 19-1, Panel B, represents these "**whole loans**"—mortgages that are not securitized and are retained in the LIC's own portfolios. Although life insurers sometimes sell these loans into pools for the Commercial Mortgage-Backed Securities (CMBS) market and may also invest in CMBS securities, their direct holdings remain significant.

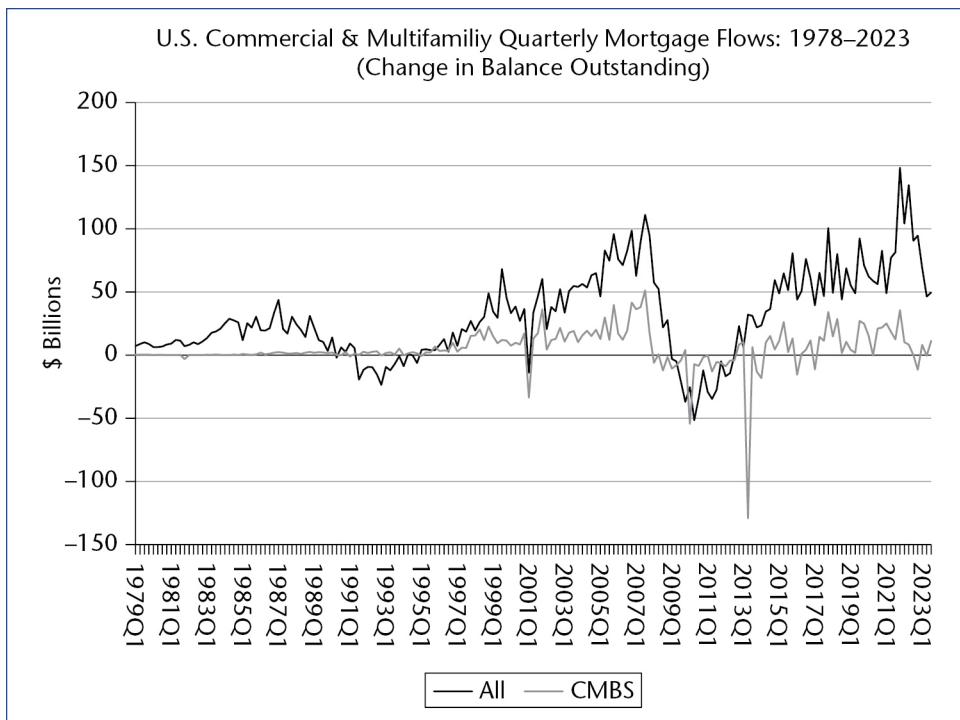


EXHIBIT 19-3 U.S. Commercial and Multifamily Quarterly Mortgage Flows, 1978–2023 (Change in Balance Outstanding).

Source: Based on data from the Federal Reserve.

The CMBS component of Panel B represents over \$960 billion of securities backed by commercial mortgage pools. CMBS lending grew rapidly from its inception in the mid-1990s and became a major source of financing for medium to large-sized commercial properties before the 2008 financial crisis, peaking at over \$800 billion in 2007–2008. This included many medium-sized “conduit loans” issued by various originators, including commercial and investment banks. The near-collapse of the CMBS industry in 2008–2009 was a major factor in the commercial real estate crash of that period. In response to the crisis, many regulatory changes were implemented, most notably through the **Dodd–Frank Wall Street Reform and Consumer Protection Act of 2010**, commonly referred to as Dodd–Frank. Dodd–Frank introduced significant reforms to the CMBS market, including the requirement for issuers to retain at least 5 percent of the credit risk in the securities they create, known as “skin in the game,” to better align their interests with those of investors. The act also mandated greater transparency in the securitization process, stricter regulatory oversight, and higher capital requirements for banks holding CMBS. While these reforms were designed to enhance the stability and resilience of the CMBS market, they also increased the cost and complexity of issuing these securities.

During the 2000s, the Government Sponsored Enterprises (GSEs, such as Fannie Mae and Freddie Mac) assumed a more significant role in providing debt capital to multifamily housing. In the aftermath of the 2008 financial crisis, they were among the few major sources of commercial mortgage credit still available. As of 2023, approximately 24 percent of multifamily development are funded by the GSEs, while 25 percent are funded by CMBS, as shown in Panel C.

Exhibit 19-3 shows the quarterly magnitude of the net flow of capital into commercial mortgages (including multifamily) from 1979 to 2023, including the share that was CMBS. The flow of mortgage capital into commercial real estate has been famously cyclical, tending to peak in the big property booms just before the big crashes. This is seen in Exhibit 19-3 during the 1980s and again in the 2000s. At the peak of the 2000s property boom, mortgage capital was flowing into

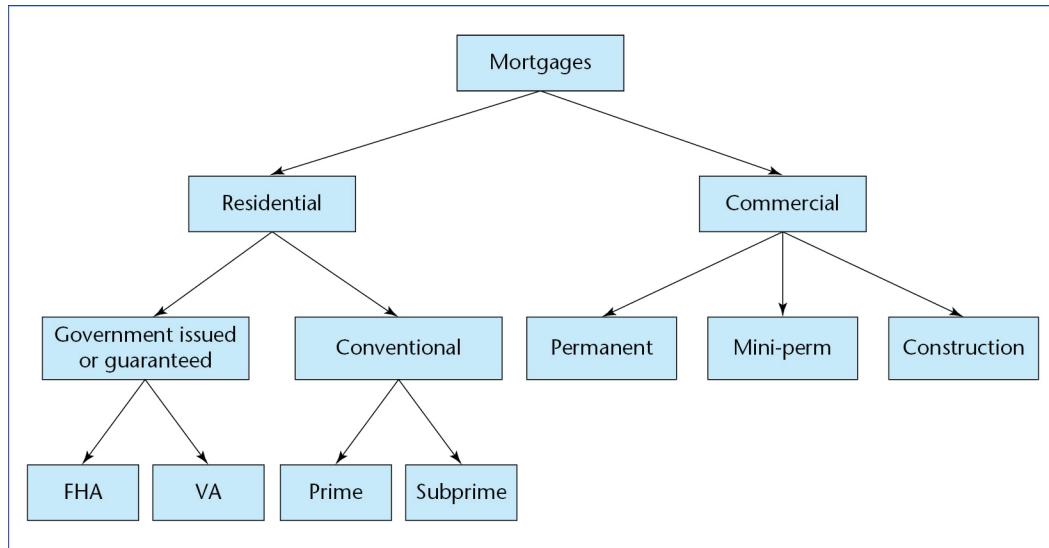


EXHIBIT 19-4 Schematic of U.S. Mortgages.

U.S. commercial real estate at an annual rate of over \$300 billion, and well over \$100 billion of that was through the CMBS industry. By that time, CMBS loans included a relatively small number of large loans on prime properties as well as a very large number of smaller- to medium-sized loans on typical commercial properties. The CMBS industry combined these very different types of loans and properties into a type of mortgage pool known as a “fusion,” because it included both types. Exhibit 19-3 graphically shows how in the crashes following the booms the flow of mortgage capital actually becomes negative in the United States, as lenders retrench and seek to reduce their real estate exposure, selling and marking down the mortgages they hold and not refinancing all the debt that is paid off. These huge flows of debt capital, in both directions, both reflect, and cause, the property pricing cycles.

The COVID-19 pandemic brought unprecedented challenges to the commercial real estate market, leading to a dramatic rise in office vacancies as businesses rapidly transitioned to remote work. In 2024 and beyond, the commercial mortgage market was poised for a significant wave of maturities, with 20 percent of the outstanding commercial mortgages set to mature, marking a substantial increase from the previous year. This surge was largely due to the extensions and modifications granted in 2023, as lenders and servicers provided flexibility amid economic uncertainty. Loan modifications and an increase in distressed real estate, especially within the office market, were driven by the higher interest rates and lower values making refinancing a challenge.

19.1 BASIC SCHEMATIC OF MORTGAGES AND OVERVIEW OF THE U.S. MORTGAGE INDUSTRY

While we can categorize mortgages in many ways, Exhibit 19-4 presents a simple schematic useful for characterizing the mortgage industry in the United States in the first quarter of the twenty-first century. Broadly speaking, mortgages are divided into **residential mortgages** and **commercial mortgages**. The former are secured by owner-occupied single-family homes, the latter by income-producing property.¹ Residential and commercial loans make up two distinct branches of

¹ Most lenders categorize loans for multifamily dwellings with up to four units as residential, especially when the borrower lives in one of the units. Loans for larger apartment buildings and complexes, however, are typically treated and classified as commercial mortgages.

the industry and are typically administered separately. These two sides of the business differ in several respects:

- Individual residential loans are generally much smaller in size but far more numerous than commercial loans.
- Residential owner-occupied properties generate no income, so the lender depends on the individual borrower's income to service the loan, while commercial loans can be serviced from the income produced by the property securing the debt.
- Residential borrowers are usually not financial or business professionals and only seek loans infrequently, typically every 5 to 10 years.² In contrast, commercial borrowers are often businesses or financial entities staffed by professionals with significantly more financial expertise than the average homeowner.
- Commercial properties tend to be more unique, while single-family homes tend to be relatively homogeneous.
- Social and political concerns influence government involvement more so for residential loans than for commercial loans. This is reflected in the distinct statutory and common laws governing foreclosure and bankruptcy for residential loans compared to commercial loans, as well as the dominance of the GSEs in the secondary mortgage market

As a result of these differences, the residential mortgage business has become a “mass production,” fairly standardized industry, while commercial mortgages, especially larger loans, remains more of a “custom shop,” where individual loans are crafted and negotiated, to some extent, one deal at a time (although this is less true in the CMBS part of the industry). Also, the federal government established an extensive regulatory oversight that has helped to standardize the residential mortgage business. Residential mortgages are generally considered part of the “consumer finance” industry, rather than business or investment finance, and are subject to different regulations under the Dodd-Frank Act. Our main focus in this chapter is on commercial mortgages.

Within each of these two main branches of the industry are further subbranches and categories of loans. For example, residential mortgages are traditionally divided into **government-insured** and **conventional loans**. Government-insured loans are backed by agencies such as the Federal Housing Administration (FHA), the Department of Veterans Affairs (VA), or the Department of Agriculture (USDA), with the federal government insuring the lender against loss in the event of default and foreclosure. This insurance allows these loans to cover larger portions of the property's value, with VA loans often covering up to 100 percent of the house value. In contrast, a conventional loan is a mortgage that is not insured or guaranteed by any government agency. Conventional loans have no government-provided default insurance. However, private mortgage insurance is often purchased, enabling loan-to-value (LTV) ratios up to 95 percent of the house value. Without private mortgage insurance, conventional loans are typically limited to no more than 80 percent LTV.³

Loan standardization and the widespread availability of default insurance, and the related efforts of the federal government starting in the 1930s and continuing into the 1970s, helped to establish an extensive and well-functioning **secondary market** for residential mortgages in the United States. A secondary market is one in which mortgages are bought and sold by third parties. The **primary market** is where the loans are **originated** and issued initially. In the primary market, the lending institution or mortgage company issuing the mortgage is the buyer of the loan, in the sense that it is providing the money up front in return for the promised future cash flows (i.e., taking the long position in the mortgage as an investor). The borrower is the seller of the loan, in the sense that it is receiving the money up front, in effect taking a short position in the mortgage as an investor. Once issued, however, the mortgage itself is a capital asset, a claim on a stream of future cash flows, and

² Although when interest rates are dropping, we will observe more rapid refinancing of existing mortgages.

³ Loans above 80 percent LTV without insurance or where the credit of the borrower is questionable are referred to as subprime loans. These include loans with less documentation, called No-doc loans or Alt-A loans, issued prior to 2008.

this asset can be traded (if properly constructed legally), that is, sold by the original issuer to a third party, who may then sell it again to another party, and so forth. This is the secondary market. The development of such a market for home mortgages was pioneered in the United States as a means to promote home building and home ownership coming out of the Great Depression and World War II. By using money derived from the bond market and providing liquidity for such investors, the secondary market has facilitated the supply of financial capital to housing and in turn helped to transform American society and the American landscape. Today, the secondary market allows lenders to sell mortgages they originate, thereby converting these long-term, illiquid assets into liquid capital. This ability to sell mortgages provides several benefits to mortgage issuers. It reduces the risk associated with holding large amounts of long-term loans on their balance sheets, freeing up capital that can be used to issue new loans or invest in other areas.

In the case of residential mortgages, the secondary market experienced extensive **securitization** beginning in the early 1970s. That is, not only was there a secondary market for individual mortgages (which had become dominated by Fannie Mae and Freddie Mac), but large numbers of individual mortgages were pooled with large numbers of small homogeneous securities (like shares of stock, or “units”) being sold into the bond market, based on the underlying mortgage pool. Holders of these securities receive shares of the cash flows received by the pool of underlying mortgages. In fact, most residential mortgages in the United States are not retained by their original issuing institution, but rather are sold into the secondary market to be securitized largely as MBS.

By contrast, several barriers inhibited the development of a secondary market for commercial mortgages for many years. Commercial properties tend to be larger and more heterogeneous than residential houses, and they do not benefit from government-support programs such as FHA insurance. The lack of standardization and default protection, combined with the smaller number and scale of the commercial mortgage industry, made securitization difficult. However, in the early 1990s, and encouraged by the successful securitization of troubled loans after the RTC (Resolution Trust Corporation) took over many troubled financial institutions, widespread securitization of commercial mortgages took off with the development of the **commercial mortgage-backed securities (CMBS)** industry, though with some notable differences from residential loan securitization. Because no default insurance exists for commercial loans, bond-rating agencies had to develop methods to classify CMBS securities by their default risk, allowing the bond market to price these securities accurately. The growth of CMBS has led to more standardization in commercial mortgages, but still not nearly as much as in residential loans. Up until the early 2000s, most conventional residential mortgages were financed by capital from Fannie Mae and Freddie Mac. These GSEs were private corporations, but chartered by the federal government, and it was widely believed that the government provided an unofficial but implicit guarantee for them. This enabled them to obtain capital at lower cost, in effect, an implicit government subsidy for housing. During the 2008 financial crisis, the government did step in to bail out these institutions, leading to their receivership and government ownership. Since then, there have been various proposals to reform Fannie Mae and Freddie Mac, ranging from winding down their operations to fully privatizing them. However, significant reform efforts have stalled in Congress, leaving the GSEs under government control as of 2024.

Although Fannie Mae and Freddie Mac’s exit from government oversight is projected around 2030 or later, they have continued to serve as the primary source of financing for the crippled housing market after the crash of 2008 including financing for multifamily housing. Because of the GSEs dominance and pioneering role in providing a secondary market for mortgages and in securitization, Fannie and Freddie long provided a de facto source of standardized underwriting guidelines and procedures for conventional residential mortgages, thereby strongly influencing how, and in particular, *how carefully*, mortgages were issued, attempting to filter out borrowers who would have trouble paying the debt.

Within the commercial mortgage industry, an important distinction exists between **permanent loans** and **construction loans**. Construction loans are made specifically for the purpose of financing a construction project, whether it’s for single-family homes or income-producing properties. These loans are expected to be repaid once the construction is completed, or soon thereafter. In contrast, permanent commercial mortgages are long-term loans designed to finance a completed, fully operational income property. The economic and financial characteristics of these two types of loans are quite different.

Construction loans are relatively short-term, typically lasting one to three years, while permanent commercial mortgages usually have terms of five to ten years, sometimes longer. In construction loans, funds are disbursed incrementally as the project progresses, often with limited or no payments from the borrower until completion, at which point the entire loan, along with any accrued interest, is due. On the other hand, permanent loans typically involve a lump sum disbursement up front, with the borrower repaying the loan with interest over time. A notable development in the mortgage market during the 2000s was the emergence of interest-only loans, used in both residential and commercial markets, where no principal is paid until after a set period, which, for commercial mortgages, could extend to the entire term of the loan. Construction loans have extensive default risk, as the underlying building does not yet exist when the loan is made. However, they have relatively little **interest rate risk**, which is the risk that the value of the loan will fall in the secondary market as a result of a rise in interest rates in the bond market. The fact that construction loans are of short duration and often made at **floating interest rates** (interest rates that move with the prevailing shorter-term interest rates in the bond market), protects construction loans from interest rate risk. On the other hand, permanent loans have less default risk (as they are typically secured by a fully operational property), but they often have considerable interest rate risk for the investor (unless they are **adjustable rate mortgages**, or **ARMs**, which is becoming more the rule for many commercial mortgage loans).

As a result of these differences, construction loans and shorter-term permanent loans are traditionally issued by commercial banks and thrift institutions, while longer-term commercial mortgages are often placed with life insurance companies and pension funds and with the GSEs (for multifamily mortgages). Banks and thrifts are **depository institutions** whose liabilities are of short duration since depositors can remove their money largely on demand. They need to match these short-term liabilities with short-term assets, such as construction loans. Life insurance companies and pension funds typically have long-term liabilities, such as life insurance policies and pension benefit obligations, which often come with a high degree of predictability in their future cash outflow requirements. To manage these obligations, these institutions need to align their long-term liabilities with stable and reliable long-term cash inflows, which permanent mortgages can effectively provide.

Because of the extensive default risk in a construction project, the construction lender must be very familiar with the local real estate space market and the local real estate developers and construction firms. Commercial banks and thrift institutions have extensive systems of local branch offices and local loan officers that can develop this kind of expertise. Permanent lenders buying whole loans are often large national and international institutions, far removed from Main Street. However, they usually work through **mortgage bankers** and **mortgage brokers** who have local expertise in the local real estate industry in order to place their mortgage capital.⁴

In the traditional functioning of the commercial mortgage industry, construction and permanent lenders often work as a team to provide construction finance. The construction lender will not commit to provide the construction loan until a permanent lender has agreed to provide a permanent loan on the completion of the project. In this case, the permanent loan is referred to as a **take-out loan**, as it is used to pay off (or “take out”) the construction loan. In the event a construction loan is made with the possibility of extension into the operational period of the property, it is called a “mini-perm.”⁵

⁴ Mortgage bankers issue loans but do not generally hold onto them as long-term investments, rather selling them immediately to long-term investors such as life insurance companies and pension funds. Mortgage brokers do not issue loans, but screen loan applicants and bring candidates to the issuing institutions.

⁵ Sometimes a third type of loan is used to bridge the construction loan and the permanent loan. This may be necessary for projects built “on spec,” that is “speculative” developments, in which few if any tenants for the building have committed themselves in advance to lease space in the completed building. For such projects, there may be an extensive lease-up period between the completion of the construction and full or nearly full occupancy of the building. The construction lender may want to be paid when the project is physically complete, but the permanent lender may not be willing to disburse funds until the building is substantially occupied. Bridge financing is sometimes referred to as a “mini-perm” loan, and it may be provided by either the permanent lender or the construction lender. Construction finance, and the analysis of the financial feasibility of development projects, was introduced in more depth in Part IV of this book.

A NOTE ON TERMINOLOGY: EE VERSUS OR

Mortgage terminology is often rather arcane. Two widely used terms that are often confusing to new students are *mortgagee* and *mortgagor*. The mortgagee is the lender, and the mortgagor is the borrower. Here's a way to help remember these: the borrower is the property *owner*. *Owner* begins with an O, so the or ending is appropriate. The same trick works for the terms *lessor* (the landlord, or property owner) and *lessee* (the tenant or renter).

19.2 BASIC LEGAL CONCEPTS AND TERMINOLOGY FOR MORTGAGES

Whether residential or commercial, permanent or construction, all mortgages are based on certain fundamental legal concepts and structures. Mortgages also typically involve a number of specific provisions and terms that may sound rather like "legalese," but you must be familiar with them in order to work with real estate debt at a practical level. In this section, we will introduce some of these basic considerations.

19.2.1 LEGAL STRUCTURE OF MORTGAGES

A mortgage is **secured debt**, which means that specified collateral can be used by the debt holder, or **mortgagee**, to obtain the funds owed if the borrower, or **mortgagor**, fails to pay what is owed under the loan. In a real estate mortgage, the collateral is real property. The mortgage is technically made up of two separate but related legal documents: a **promissory note** and a **mortgage deed**.

The promissory note establishes the debt. It is a written, signed contract between the borrower and lender. In this contract, one legal person (the mortgagor) promises to pay another legal person (the mortgagee) the cash flow amounts specified in the loan, in return for the loan.

The mortgage deed, also known sometimes as a **security deed** or a **deed of trust**, secures the debt by conveying, or potentially conveying, the ownership of the collateral from the borrower to the lender. In essence, if the borrower fails to live up to the terms and conditions specified in the promissory note, then the mortgage deed enables the lender to acquire the collateral property for the purpose of obtaining what is owed. The mortgage deed should be recorded in the title recording office appropriate to the specified property. The date it is recorded will, in the absence of contravening factors, establish the priority that the lender has in his claim to the collateral property.

There are two types of legal bases for the lender's claim on the collateral property. In some states, mortgages are governed by what is called the **lien theory**: the mortgagee holds a lien on the collateral property. This gives her the right to take the property to force a foreclosure sale in the event the borrower fails to perform under the promissory note, but it does not give the lender the title to the property. In contrast, in some states, mortgages are governed by what is called the **title theory**.⁶ Under this theory, the mortgagee holds the ownership title to the collateral property until the borrower is released from the promissory note commitment (normally by the loan being paid off). The borrower technically only retains the right of use and possession of the property, as long as the provisions of the promissory note are kept. In title theory states, the lender can generally take possession of the property more quickly and easily in the event of a default.

⁶ The title theory states are Alaska, Georgia, Nebraska, South Dakota, Utah, Arizona, Idaho, Nevada, Tennessee, Virginia, Colorado, Mississippi, North Carolina, Texas, West Virginia, Missouri, Oregon, Washington, Wyoming, and District of Columbia. There exists some debate as to whether California is title theory or lien theory state. While many mortgage law experts consider it a title theory state, Section 5 of the California Department of Real Estate's reference book states, "It is settled law that California is a 'lien' and not a 'legal title' theory state when imposing encumbrances/liens against the title of real property." It is listed here for the sake of inclusion.

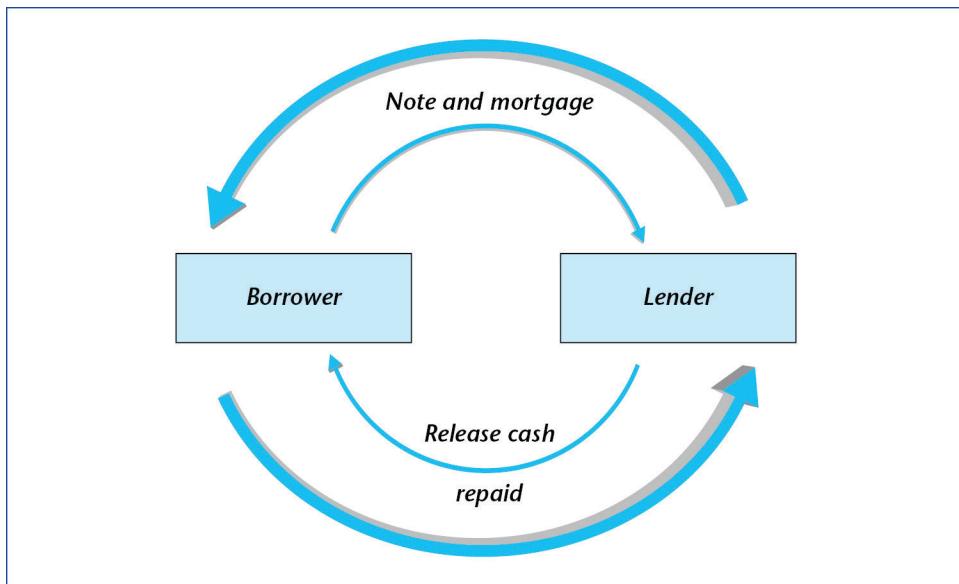


EXHIBIT 19-5 Mortgage Deed Relationships.

The essential relationship between the borrower and lender regarding the note and deed is portrayed in Exhibit 19-5. In return for cash up front the borrower gives the lender a note and deed. When the loan is paid off, the lender releases the borrower from the note and returns the deed.

In some states, the legal structure of the mortgage is more complicated, based on a deed of trust, as depicted in Exhibit 19-6. In this case, a third party known as the trustee is involved in the mortgage, in addition to the borrower and lender. When the money is loaned, the borrower gives the ownership title of the property to the trustee, who holds the title for the benefit of the lender during the time the loan is in effect. If the loan is paid off, the trustee returns the title to the borrower. If the borrower defaults, the trustee has the power to sell the property on behalf of the lender, using the proceeds of the sale to compensate the lender for what is owed under the loan. In this case, the trustee has what is called the **power of sale**. In some states, the mortgagee has the power of sale even without the deed of trust arrangement. In other states, only a court can actually sell the collateral property in foreclosure, in what is called a **judicial sale**.⁷ In general, the foreclosure sale process tends to be quicker and less administratively costly in power-of-sale states.

19.2.2 PRIORITY OF CLAIMS IN FORECLOSURE

As noted, if the borrower defaults on his obligations under the promissory note, the lender can force a **foreclosure** sale of the collateral property and use the proceeds from that sale to obtain the debt owed. However, there may be multiple claimants to the proceeds from the foreclosure sale. Anyone with a lien on the collateral property has a claim. There may be more than one mortgage on the property, and there may be other liens besides those held by mortgage lenders. For example, the local property tax authority can place a lien on a property if the owner has not paid the property taxes due.

⁷ The major states in which judicial sales predominate for commercial mortgages are Connecticut, Delaware, District of Columbia (sometimes), Florida, Hawaii, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana (executory proceeding), Maine, Nebraska (sometimes), New Jersey, New Mexico, New York, North Dakota, Ohio, Oklahoma (if the homeowner requests it), Pennsylvania, South Carolina, South Dakota (if the homeowner requests it), Vermont, and Wisconsin.

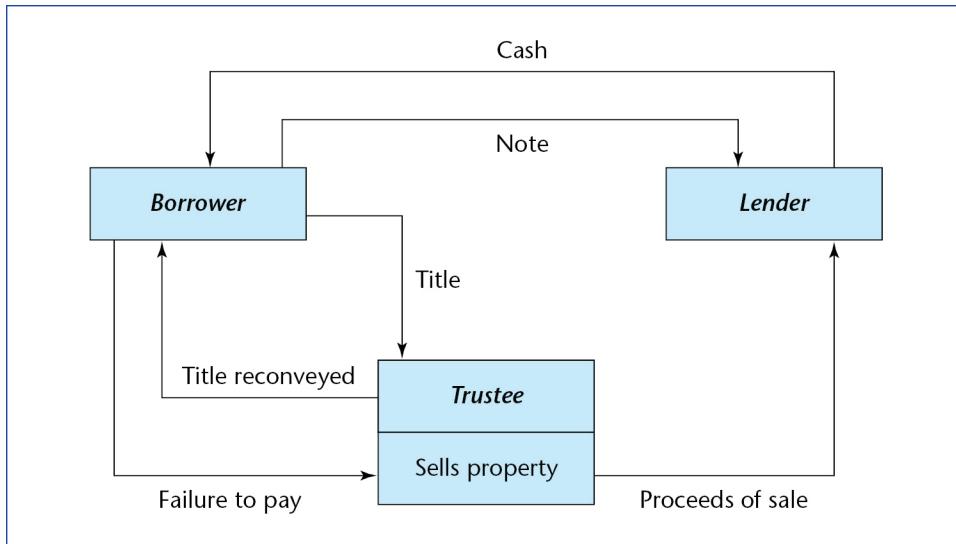


EXHIBIT 19-6 Deed of Trust Relationships.

Additionally, a construction firm that has performed work on the property can place a mechanic's lien if it was not paid in full for its services.

When property has multiple lien holders, the claimants are arranged in a "pecking order" of priority that is generally established by the date of recording of the lien, with the following exceptions:

- Property tax liens come first.
- Sometimes mechanics' liens can take priority over others.
- Explicit subordination clauses can reduce priority.
- Bankruptcy proceedings may modify debt holder rights.

Property tax liens are always given first position, regardless of when they were recorded. Statutory provisions often allow construction firms a specific window of time to record mechanics' liens after work has been performed and not paid for. These liens can take precedence over other liens that were recorded before the end of this period. A loan may include a subordination clause, which explicitly makes it secondary to other loans issued later. This is common in seller-provided mortgages. Lastly, if the borrower declares bankruptcy, even secured creditors may have their claims modified by the bankruptcy court, especially in the case of commercial mortgages.

Apart from these exceptions, the priority of claims in foreclosure goes according to the date of the recording of the lien, which would normally be the date when the mortgage was issued. The first mortgage is the most **senior debt** in position, and would normally be the outstanding loan that was made first. If the property owner subsequently took out a second mortgage, it would be subordinate to the first mortgage, an example of **junior debt**. There may be additional junior mortgages, such as a third mortgage, a fourth, and so on. Senior claims must be fulfilled completely before any remaining available proceeds from the foreclosure sale are applied to the next highest claimant in the pecking order.

For example, imagine a scenario where \$9 million is owed on a first mortgage, \$2 million on a second mortgage, and \$1 million on a third mortgage. If the collateral property is sold in foreclosure for \$10 million, the first mortgagee will be fully paid the \$9 million owed. The second mortgagee will receive the remaining \$1 million from the sale, covering only half of the amount owed. The

third mortgagee will receive nothing, as all proceeds have been exhausted by the first and second mortgages.

It is important to note that this pecking order prevails regardless of which lien holder initiates the foreclosure suit. Subordinate lien holders will have their liens extinguished by the foreclosure suit, provided they are included in the suit and properly notified in advance. For this reason, junior lien holders who stand to gain or lose the most in the foreclosure process are most likely to actually bring the foreclosure suit. This is described by the dictum: *Redeem up, foreclose down*.

Take the previous example as a case in point. If it is known that the current market value of the property is approximately \$10 million, then the foreclosure suit would most likely be initiated by the second mortgage holder, as they are on the margin regarding whether they will obtain their full principal or not. Lien holders can bid for the property in the foreclosure sale and use their claim as payment. For example, the second mortgagee could bid \$11 million for the property and use their claim to pay \$2 million of this. They would then have to pay the first mortgagee the full \$9 million owed on the first mortgage, but the second mortgagee has now gained control of the property and can ensure it is sold for the highest possible price. The first lien is extinguished by the suit and the payment of the \$9 million, which represents the “*redeem up*” aspect of the dictum. The third lien is extinguished simply because the third lien holder was included in the suit and properly notified in advance, illustrating the “*foreclose down*” part of the dictum. If the second mortgagee sells the property for \$10 million, they would be left with \$1 million, which is half of what was owed under their loan.

The lien holder bringing the suit must identify any junior lien holders and include them in the foreclosure suit. The suit will not extinguish any validly recorded liens not included in the suit. The purchaser of the property in the foreclosure sale could end up owning the property subject to a lien that had been junior and now is senior — a danger of which any astute bidder would be aware.

19.2.3 TYPICAL MORTGAGE COVENANTS

The deed and note in a mortgage will typically include various clauses and **covenants** that outline the agreement between the borrower and lender at the time the mortgage is issued. The rights and duties of both parties should be specified clearly in these documents. Over the years, certain clauses and covenants have become common in many mortgages. Familiarity with these terms will enhance your practical understanding of how mortgages operate.

Promise to Pay This clause is present in every mortgage and means just what it says. The borrower promises to pay the lender the principal, interest, and penalties specified in the promissory note.

Order of Application of Payments This clause establishes the order in which any payments received from the borrower will be applied to different components of the debt. Typically, the standard order is as follows:

First: Expenses

Second: Penalties

Third: Interest

Fourth: Principal

If the lender has incurred any expenses that the borrower is obligated to reimburse, the borrower's payments will first be used to cover these costs. Any remaining funds will then be applied to pay off any penalties the borrower has incurred under the terms of the debt. After penalties are addressed, the payments will be used to cover the interest owed. Finally, any remaining amount from the payment will be applied to reduce the outstanding principal balance of the loan.

Good Repair Clause This clause requires the borrower to maintain the property in reasonably good condition. It is usually paired with another clause: the lender's right to inspect, which allows the

lender to enter and inspect the property periodically (typically with prior notice and at the borrower's convenience) to ensure that the property is being properly maintained. These clauses help the lender to make sure that the value of their collateral is preserved.

Joint and Several Liability This clause stipulates that when there is more than one signatory to the loan, each individual who signs is fully responsible for the entire debt.

Acceleration Clause In general, an **acceleration clause** allows the lender to "accelerate" the loan, meaning that the entire outstanding principal balance becomes due and payable immediately under certain conditions. Acceleration clauses may serve several purposes within a mortgage agreement. One common use is in the event of default by the borrower, such as when the borrower fails to make the required payments. This allows the lender to recover the entire remaining loan balance through the foreclosure sale process, rather than just the delinquent payments, penalties, and expenses. Another typical use of the acceleration clause is to enforce a due-on-sale provision.

Due-on-Sale Clause This enables the lender to accelerate the loan whenever the borrower sells the property.⁸ This is particularly valuable to the lender if market interest rates have risen since the time the loan was issued, enabling the lender to reinvest at a higher yield the capital that had been tied up in the loan. In any case, without such a clause, the lender could find themselves with a borrower who is responsible for the note but no longer owns the property securing the loan. This could cause a loss of incentive for the borrower to repay the loan. The due-on-sale clause does not require the lender to accelerate the loan on the sale of the property; it merely gives the lender that right. There may be times when the lender would be pleased to have the loan continue in effect after a sale, perhaps with the new property owner taking over the loan.

The process of transferring a loan from the original borrower to a new borrower is known as **loan assumption**.⁹ In the absence of a due-on-sale clause, the lender will not generally be able to effectively prevent the assumption of the loan by a new buyer of the property.¹⁰ Such an assumption can be valuable for both the original borrower and the new buyer of the property if interest rates have risen since the time the loan was issued. In effect, it enables the purchaser of the property to obtain a below-market-interest-rate loan, which should enable the seller of the property to sell it for a higher-than-market price. Most conventional residential mortgages in the United States have due-on-sale clauses in them, but FHA and VA loans do not, and many commercial mortgages do not.

Borrower's Right to Reinstate This clause allows the borrower to stop the loan acceleration triggered by default, provided that the borrower has cured the default (by paying all outstanding amounts) before a certain deadline, typically just before the foreclosure sale.

Lender in Possession Clause This clause gives the lender the automatic right of possession of the property in the event the borrower defaults on the loan. This enables the lender to control the leasing and maintenance of the building prior to completion of the foreclosure process. Sometimes this can help the lender to mitigate the "running down" or loss in value of the property during a prolonged foreclosure process.

Release Clause In all loans, the borrower is released from the debt, and the lender is required to return the mortgage deed and extinguish the lien on the property when the loan is paid off according

⁸ Usually, the due-on-sale clause is written such that any transfer of a substantial beneficial interest in the property triggers the acceleration provision. Thus, the due-on-sale requirement cannot be avoided by use of land contracts or other such procedures in which the title does not initially change hands.

⁹ One situation in which assignment is useful is when the original borrower is having trouble servicing the debt, not due to problems with the collateral property, but to problems from other sources. By selling the property and assigning the loan to the new owner, the health of the debt can be restored.

¹⁰ Even without a formal assumption of the mortgage, the borrower (seller of the property) will often be able to keep the original loan in effect through the use of a "wraparound" second mortgage. The seller issues a second mortgage to the buyer for an amount that covers the payments on the first mortgage.

to the terms of the promissory note, a process known as defeasance or reconveyance. In some cases, additional specific release clauses may be included, such as clauses that free specific parts or amounts of the original collateral property upon repayment of portions of the debt principal. This is particularly useful when the original collateral is intended to be sold off gradually in parts or parcels, as is often the case with housing tract developments and subdivisions. In another use, an original borrower may be released from the debt when the mortgage is assumed by another debtor.

Prepayment Clause This provision grants the borrower the right, but not the obligation, to pay off the loan before its maturity.¹¹ This can be valuable for the borrower when interest rates fall below the contract rate on the loan, as it enables the borrower to refinance the mortgage at lower interest rates by paying off the original mortgage with proceeds from a new, lower-interest-rate loan. Lenders are aware of this possibility when they issue the mortgage and typically charge an interest rate high enough to compensate for the value of this option provided to the borrower. Unless a prepayment clause is explicitly stated in the mortgage agreement, the borrower does not have the right to pay the loan off early, that is, prior to its originally stated maturity. Virtually all residential mortgages in the United States have prepayment clauses.¹² Many commercial mortgages do not have prepayment clauses, or have such clauses that only become effective after a period of several years, often referred to as a “lockout” or as the borrower being ‘locked in’ to the loan. Often, commercial loans technically permit prepayment but set prepayment penalties so high as to effectively eliminate the value for the borrower of refinancing the loan at a lower rate.¹³ This eliminates **prepayment risk** for the lender, which is the risk that the loan will be paid off at the borrower’s discretion before the loan’s contractual maturity. This is an important consideration because many debt investors want to use their investments precisely for the purpose of matching the maturity of liabilities with that of their assets, and so they buy debt assets with a view toward holding them to maturity.¹⁴ Sometimes prepayment penalties kick in after a lockout period and then diminish with a sliding scale approaching the loan maturity.

During the 2000s, the evolution in the CMBS industry was toward conduit loans having what is called **defeasance** provisions. This provided that the borrower could pay off the mortgage but only by providing setting up a series of (typically) U.S. Treasury bond securities that would assure the mortgage holder (the CMBS trust) receives all of the payments that the borrower would otherwise have made under the loan. This actually improved the situation for the mortgage holder, because they received the same cash flows as they otherwise would have, only now with the credit quality of U.S. government bonds. Obviously, this process was costly to the borrower, and so would only be worthwhile if the borrower faced a strong incentive to get out of the mortgage obligation.¹⁵

¹¹ This is broadly the same type of option that exists in callable corporate bonds. In effect, it gives the borrower a *call option on a bond*, in which the bond has cash flows equivalent to the remaining cash flows on the callable debt, and the exercise price of the option is the outstanding loan balance (plus prepayment penalties, i.e., what one would have to pay to retire the debt).

¹² Government regulations require that adjustable-rate residential mortgages permit prepayment without penalty. Most fixed rate residential loans also allow prepayment with little or no penalty as of 2024.

¹³ This still gives the borrower some flexibility but removes the interest rate risk for the lender. For example, the prepayment penalty may be specified to be sufficient so that the lender’s original contract yield is maintained over the original contract maturity of the debt, called yield maintenance, even though the lender would have to reinvest the pre-maturely prepaid principal at market interest rates below that contractual yield for the remainder of the original loan’s lifetime. Yield maintenance based on the present value of the remaining loan term is often based on the presumption that the paid-off loan proceeds are reinvested in Treasury bonds of a specified duration so that the lender assumes zero yield risk when prepayment occurs.

¹⁴ If they can do this, they also avoid exposure to interest rate risk in the sense that they would have to reinvest their capital at lower interest rates prior to the end of their investment target horizon.

¹⁵ In the 2000s the incentive was not usually merely to obtain a lower interest rate mortgage, but rather to obtain a *larger* mortgage, thereby “cashing out” some of the increased equity in the property resulting from the rapidly growing property values during the unprecedented commercial property price boom of the 2004–2007 period. (Or the defeasance might have been associated with a property sale that was similarly motivated by the run-up in prices.) In many cases the properties that were “defeased” during the peak of the 2000s boom subsequently became “distressed assets” after the crash of 2008–2009.

Subordination Clause As previously noted, this provision makes the loan subordinate to other loans that the borrower may obtain, meaning it has a lower claim priority for the lender in the event of foreclosure, even if those other loans were recorded after the loan in question.¹⁶ Subordination clauses are often used in seller loans and subsidized financing to enable the recipient of such financing to still obtain a regular first mortgage from normal commercial sources.

Lender's Right to Notice This is a provision in junior loans (i.e., subordinate loans, such as second mortgages) requiring the borrower to notify the lender if a foreclosure action is being brought against the borrower by any other lienholder. Junior lien holders might choose to assist in resolving the borrower's default on the other loan or seek an alternative solution to avoid foreclosure. This is because junior lien holders will stand to lose much more in the foreclosure process than senior lien holders, due to their being lower in the pecking order for the foreclosure sale proceeds.¹⁷

Future Advances Clause This provides for some or all of the contracted principal of the loan to be disbursed to the borrower at future points in time subsequent to the establishment (and recording) of the loan. This is common in construction loans, in which the cash is disbursed as the project is built. Even though some of the loan principal is not disbursed until later, the priority of the lien applying to all of the principal is established based on the time of the initial recording of the mortgage.

Exculpatory Clause This removes the borrower from responsibility for the debt, giving the lender no recourse beyond taking possession of the collateral that secures the loan. Without an exculpatory clause, the lender can obtain a deficiency judgment and sue the borrower for any remaining debt owed after the foreclosure sale (i.e., in the event the foreclosure sale proceeds are insufficient to recompense the lender for all that is owed). Loans containing exculpatory clauses are known as **nonrecourse loans**. They are common in the commercial mortgage business in the United States, especially for borrowers that have a solid track record of past performance.¹⁸

19.3 DEFAULT, FORECLOSURE, AND BANKRUPTCY

Although most mortgages are repaid in full without issues, problems are not uncommon. On average, over 15 percent of long-term commercial mortgages in the United States end up facing serious default, and a foreclosure, bankruptcy, or workout situation, at some point before the maturity of the loan.¹⁹ Technically, a **default** occurs whenever the borrower violates any clause or covenant in

¹⁶ Recall that claim priority is normally established by the date of recording of the mortgage deed. The subordination clause overrules this priority.

¹⁷ This provision would enable the junior mortgagee to find out about the foreclosure action earlier than would necessarily occur when she is officially joined in the foreclosure suit. Recall our example in the previous section in which there were three lien holders. The third lien holder might want to cure the borrower's default on the other loans to prevent the foreclosure sale in which the third lien holder stands to lose his entire claim. Perhaps the property market will turn around or the property itself can be turned around so that it would be worth more than \$12 million, instead of the current expected value of \$10 million. At \$12 million, there would be enough value to make whole even the third lender.

¹⁸ Even in the absence of an explicit exculpatory clause, commercial mortgages can often be made effectively nonrecourse by use of "single-asset" borrowing entities that have limited liability, such as a corporation that owns nothing other than the property that is being financed by the mortgage. However, astute lenders wanting to avoid a nonrecourse loan situation can require the parent holding company or an individual with large net worth to sign the note with joint and several liability, or they may require additional collateral for the loan. On the other hand, an opposite situation can occur. The lender may want protection against possible borrower bankruptcy and loan modification "cramdown" during a Chapter 11 bankruptcy proceeding (see next section). A "bankruptcy-remote" single-asset entity is set up to borrow the loan. However, as CMBS lenders to General Growth Properties found out after 2009, care must be taken that the bankruptcy-remote borrowing entity really is sufficiently independent of any larger corporation such that the larger corporation's bankruptcy does not jeopardize the mortgagee's claim on the property.

¹⁹ Data on permanent loans issued by life insurance companies between 1972 and 1992 indicate a lifetime cumulative default rate roughly between 10 percent and 25 percent, depending on the nature of the property asset market in the years subsequent to the issuance of the loan. Default is most likely during the third through seventh year of the loan's lifetime. See Esaki, l'Heureux, and Snyderman (1999). Their working definition of a serious default is payment delinquency greater than 90 days. The incidence and cost of commercial mortgage foreclosure will be discussed in more depth in Chapter 18.

the mortgage agreement. However, many violations are not serious in that they do not pose a threat to the lender's yield-to-maturity and are ignored by the lender. Serious defaults generally have to do with failure to make the stipulated loan payments.

When a borrower defaults, lenders have several possible courses of action. Broadly speaking, these can be categorized as being either **litigious** or **nonlitigious** in nature. Litigious actions involve the courts and formal legal proceedings, which can ultimately lead to foreclosure. When lenders choose to delay or avoid this approach, they are exercising **forbearance**. Forbearance became particularly relevant during the COVID-19 pandemic, when the U.S. government implemented a mortgage forbearance program as part of the Coronavirus Aid, Relief, and Economic Security (CARES) Act in March 2020, allowing borrowers to temporarily pause or reduce their mortgage payments. Under the CARES Act, homeowners with mortgages backed by federal entities such as Fannie Mae, Freddie Mac, the FHA, and the VA were entitled to request forbearance for up to 180 days, with the option to extend for an additional 180 days. While some commercial lenders offered forbearance or loan modifications on a case-by-case basis, there was no government-mandated forbearance program for commercial mortgages. Nonetheless, some commercial borrowers successfully negotiated temporary relief directly with their lenders, depending on the terms of their loan agreements and the lenders' willingness to accommodate such requests during the pandemic.

As a general rule, lenders prefer to resolve problems without recourse to litigious actions, if possible, thereby saving legal expenses and retaining more flexibility to deal with the borrower in a less formal, less adversarial, and often more expeditious manner.

When nonlitigious options are insufficient, lenders face several types of judicial actions they can take through the courts. As a first step, they can *sue for specific performance* under the promissory note. This involves getting a court to officially require the borrower to perform some specific act (such as fix a default). This is often used as a "shot across the bow," warning the borrower how seriously the lender views the default.²⁰ A second level of recourse is for the lender to sue for damages under the promissory note, without invoking the mortgage deed. This is less expensive and administratively involved than the third step, which is to invoke the mortgage deed to bring a foreclosure action.

The foreclosure process allows a forced sale of the collateral property, with the proceeds being used to compensate secured creditors for what is owed them, as much as possible. Any remaining proceeds from the foreclosure sale, after payment of debts and expenses, go to the borrower, but there is normally nothing left over by that point. Foreclosure is generally regarded by both lenders and borrowers as a last resort because it is expensive and slow and represents a rather public type of failure of the loan for both parties.

If the foreclosure sale does not provide funds sufficient to pay the secured debt holders all that they are owed, the holders of remaining nonextinguished debts can obtain a **deficiency judgment**. This allows the lender to sue the borrower for the remainder of the debt, provided the mortgage did not contain an exculpatory clause (as previously described). However, in practice, deficiency judgments are often of limited value, as the borrower may by this time have little other assets or income that can be attached, or the borrower may declare **bankruptcy**.

Bankruptcy on the part of the borrower is in fact a major danger to commercial mortgage lenders prior to foreclosure, even if the collateral property securing the loan contains sufficient value and income-generating ability to service the debt. The most serious threat to mortgage holders when a commercial loan borrower files for Chapter 11 bankruptcy is the possibility of a **cramdown**, a situation that a bankruptcy court allows a borrower to modify the terms of a secured loan without the consent of the lender, effectively "cramming down" the loan to a lower value. This often occurs when the value of the collateral securing the mortgage is less than the amount

²⁰ To some extent, the litigious and nonlitigious tracks can be pursued in parallel, in a kind of "good cop/bad cop" routine.

owed.²¹ The idea behind Chapter 11 is to allow a potentially salvageable business to continue operating while it works itself out from under excessive debt. Under Chapter 11, all of the borrower's creditors are lumped together, and they all may be forced to accept a restructuring of their debts, including secured lenders whose collateral might be sufficient to service their loans by themselves. In effect, healthy properties could be used to bail out sick properties or businesses at the expense of the holders of mortgages on the healthy properties. In effect, a restructuring of the mortgage is forced by the bankruptcy proceeding onto the lender (crammed down the lender's throat, as it were). Because of this, the mere possibility or threat of bankruptcy can be a weapon used by borrowers against lenders.²²

19.4 NONLITIGIOUS ACTIONS AND WORKOUTS

As noted, lenders generally view foreclosure as a last resort and prefer to address loan problems before they get to that point (and before a Chapter 11 declaration can cause them to lose negotiating leverage). If applied in a timely and deft manner, nonlitigious actions can often obtain the most successful resolution of loan problems.

19.4.1 COSTS OF FORECLOSURE

The value of nonlitigious solutions to debt problems may be viewed as the avoidance of the costs of foreclosure. Foreclosure involves a variety of costs, nearly all of which are **deadweight burdens**, that is, a loss to one or the other side of the loan without a commensurate gain on the other side. In Chapter 15, we noted that as a part of the "costs of financial distress" (COFD), even the possibility of such costs can reduce the value of levered equity in property ownership, making excessive debt undesirable from the borrower's perspective. Foreclosure costs can be significant, often representing a substantial portion of the property's value and the remaining loan balance. These costs typically include the following:

- Third-party costs, including legal and administrative expenses such as court costs, typically amount to around 10 percent of the loan balance but can be significantly higher.
- Deterioration of the property during the foreclosure process is common. The borrower has little incentive to maintain the property well during this period, and lenders and courts lack property management skills. The problem with commercial property is not just the physical maintenance and repair, but also tenant relations and leasing.
- Revenue may be lost by the property, and interest payments lost to the lender, while capital is still tied up in the loan, during the time taken by the foreclosure process, which is typically close to a year in length, sometimes more.²³
- Both the borrower and the lender can suffer negative reputation effects from the foreclosure process, which is a very public procedure.

²¹ A declaration of bankruptcy under Chapter 7 is of less concern to secured lenders, as this involves complete liquidation of all of the borrower's assets, with secured debt holders obtaining the proceeds of any sale of collateral assets in much the same manner as under the foreclosure process.

²² Cramdown is of sufficiently widespread concern among commercial mortgage lenders that an industry exists providing insurance specifically against the possibility of cramdown. In some cases, lenders require borrowers to pay for this type of insurance before they will commit to the loan, or as part of a workout agreement restructuring a loan. As noted, another technique that is applicable to loans on individual properties held by large multiproperty borrowers is the formation of a "bankruptcy-remote" borrowing entity for each loan on each property.

²³ Studies by Brian Ciochetti (1996) of foreclosures on permanent commercial mortgages issued by insurance companies indicate an average length of the formal foreclosure proceeding of 9.1 months, after an average of 3.5 months of payment delinquency, for a total of 12.6 months.

- Lenders will usually have to write down the value of the assets on their books as a result of a foreclosure, whereas book values may be maintained to some extent, at least for a while, if foreclosure can be avoided. Writing down assets can cause lenders problems with reserve requirements and other regulations.

The magnitude of these and other costs can easily exceed 25 percent to 35 percent of the outstanding loan balance in total.²⁴

19.4.2 NONLITIGIOUS ACTIONS

To avoid such costs, lenders often resort to nonlitigious actions when dealing with problem loans. One option is transferring the loan to a new borrower, which is particularly common when the issue lies with the original borrower rather than the property itself—perhaps due to difficulties in the borrower's other holdings or business ventures. Another approach is negotiating a “short sale,” where the lender allows the borrower, who is at risk of default, to sell the property independently. Although the sale must be approved by the lender, it is generally accepted that the sale price will be below the mortgage balance. The lender, in effect, permits the seller to do its own marketing under the presumption that the selling price will be higher than if the property were sold as an REO (“real estate owned”) property directly from the lender. Short sales became very common in the residential market after 2008 and fairly common in the commercial industry as well in 2009–2011. Another common type of nonlitigious resolution to a mortgage problem is the procedure known as **deed in lieu of foreclosure**. In this scenario, the borrower transfers the property to the lender in exchange for being released from the debt. This outcome is nearly identical to a foreclosure but avoids the legal expenses and delays associated with it, benefiting the lender who would otherwise bear those costs. The borrower also gains an advantage, as a deed-in-lieu is less public and less damaging to their reputation than a formal foreclosure. The deed-in-lieu is a useful device when both sides agree that nothing can be done to save the loan.

A third approach is for the lender and borrower to work together to **restructure** the loan, to bring new equity partners into the deal, or otherwise creatively “work out” the problem. Indeed, this process is sufficiently common that the verb has become a noun: it is referred to as a **workout** of the loan.

19.4.3 WORKOUTS

Workouts occur in infinite varieties, with varying degrees of good or bad feelings among the parties, and with the possibility of litigious action always waiting in the wings. The basic ingredients are typically some combination of a rescheduling or forgiveness of some debt, often in return for the lender obtaining either some equity participation in the property or a greater yield in the long run. For instance, the loan term may be extended, or the borrower may be granted a temporary “holiday” from repayments. Sometimes a nonrecourse loan may be restructured with more recourse or use of additional capital, or new equity or debt partners may be brought into the deal. This may involve an assumption of all or part of the debt by the new partner. Some firms and individuals specialize in helping with mortgage workouts, acting as advisors, brokers, and mediators to the principal parties. There is even scope for entrepreneurial profit in the workout process, provided the underlying property has good long-term potential.

²⁴ The Esaki-l'Heureux-Snyderman study (1999) found an average loss severity of almost 38 percent as a fraction of the outstanding loan balance, among loans that were liquidated in foreclosure. This figure includes lost interest as well as principal and expenses.

An important consideration for the lender in mortgage restructurings is the presence of other lien holders on the property. If the subject loan is modified sufficiently, it may be deemed by the courts to be a new loan and may thereby lose its seniority among the other preexisting liens. Also, courts may invalidate a restructuring of one loan that excessively burdens other lienholders on the same property. The cooperation of all lienholders is usually necessary in a major mortgage workout.

19.5 LIMITED LIABILITY AND FORECLOSURE COSTS: THE BORROWER'S PUT OPTION AND STRATEGIC DEFAULT

Now that you have a basic understanding of how mortgages work, let's explore the investment aspect of mortgages through the lens of financial economics. Obviously, the basic idea in a mortgage is to provide the investor with a reliable cash flow stream, and this is normally what happens—a bit boring perhaps, but useful. When borrowers or loans get into trouble, however, the nature of the investment can change drastically. Investors are aware of this as a possibility (with some probability of occurrence) long before it occurs in reality. Two features of the typical commercial mortgage are of particular importance in this regard and can interact to cause additional ramifications. We are speaking of **limited liability** on the part of the borrower and the existence of significant costs associated with the foreclosure process.

19.5.1 LIMITED LIABILITY AND THE BORROWER'S PUT OPTION

Many commercial mortgages are nonrecourse loans. In effect, the only thing the lender can take in the event of default and foreclosure is the property collateral securing the debt. The fact that this property may be worth less than the outstanding balance on the loan gives the borrower a type of option known as the **borrower's put option**. Options always have a positive value to their holders, so this option has some value to the borrower. However, this value is taken from the lender.

A put option gives its holder the right without obligation to sell a specified underlying asset at a specified price. In the context of a nonrecourse mortgage, the underlying asset is the collateral property, and the borrower's ability to default on the loan effectively gives him the ability to sell the property to the lender at a price equal to the outstanding loan balance.²⁵ By defaulting on the loan, the borrower is said to "put" the property to the lender, thereby ridding himself of a liability equal, at least in book value, to the outstanding loan balance. The lender loses an asset with this same book value. Of course, the borrower loses, and the lender gains, the value of the property.

This situation is depicted in Exhibit 19-7 as of the time of maturity of a \$100 million mortgage. The horizontal axis represents the value of the underlying property securing the loan. The vertical axis represents the value of what the lender obtains. If there were no limited liability, the lender would receive \$100 million no matter what the value of the property, as depicted in Exhibit 19-7, Panel A. With the existence of limited liability (as with the nonrecourse loan), the lender faces the situation depicted in Exhibit 19-7, Panel B. As long as the property is worth more than the outstanding loan balance, the lender will receive the whole amount due (\$100 million). If the property is worth less than the loan balance, the borrower would maximize her own wealth by defaulting on the loan and returning the property to the lender. The lender then receives only the value of the property, which is less than the loan balance. For example, if the property were worth \$90 million, the lender would receive value that is \$10 million less than the \$100 million owed. If the property were worth only \$80 million, then the gap would be \$20 million.

²⁵ To be more precise, the exercise price equals the current market value of a default-free version of the mortgage. This may differ from the outstanding loan balance ("par value" of the loan) if market interest rates have changed since the loan was issued. To simplify the basic point here, however, we can ignore this interest rate effect.

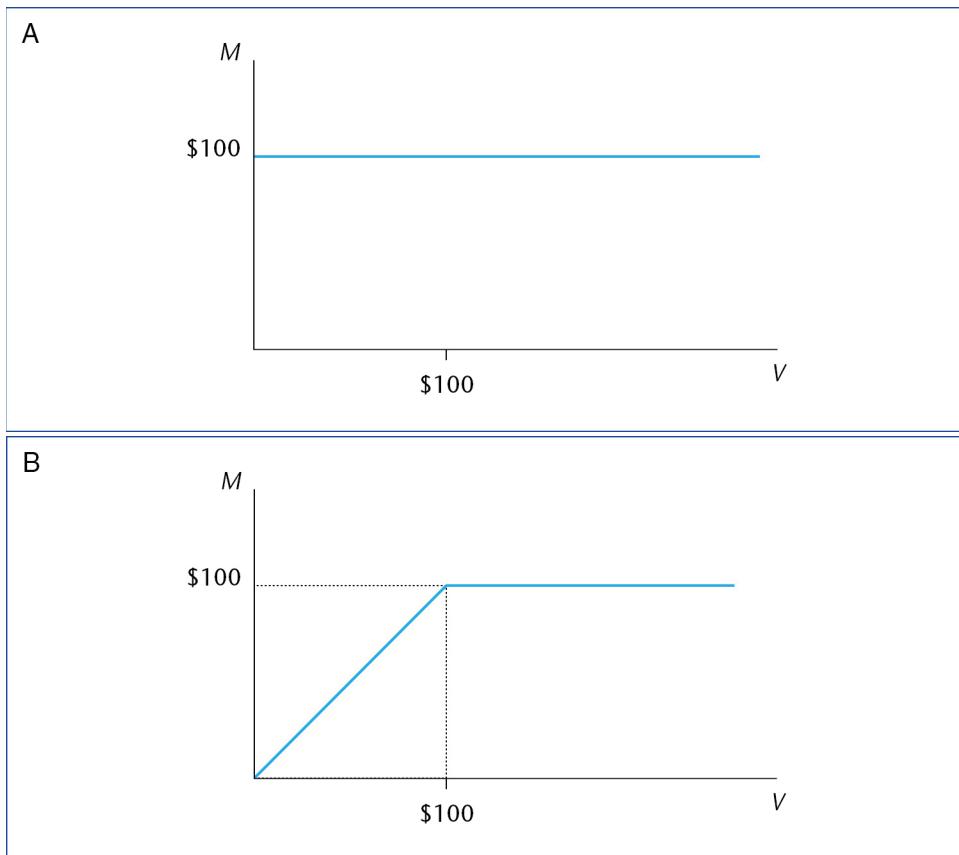


EXHIBIT 19-7 Value (M) of a \$100 Mortgage at Maturity as a Function of the Underlying Property Value (V).

The difference between the loan value to the lender without limited liability (as in Exhibit 19-7, Panel A) and its value with limited liability (Exhibit 19-7, Panel B) is the value of the put option to the borrower. To the lender, this is obviously a negative value, or cost. Lenders are well aware of this possibility when issuing loans, and mortgage investors recognize the borrower's put option when making investments. Mortgage prices and the resulting yields in the mortgage market reflect the presence of this put option. This is why mortgage yields are higher than equivalent-maturity U.S. government bonds, which do not carry default risk. On a probabilistic basis, the greater the expected losses from the borrower exercising the put option (i.e., the greater the **credit risk** or expected credit losses in the loan), the higher the contractual yield that the mortgage lender or investor will demand.²⁶

19.5.2 FORECLOSURE COSTS AND STRATEGIC DEFAULT

In the previous example, we assumed there were no foreclosure costs. If the property was worth \$90 million when the borrower defaulted, the lender would receive \$90 million in value on a

²⁶ Chapters 12 and 21 discuss in more depth the impact of credit risk on commercial mortgage interest rates and expected returns. Chapter 16 presents a simple binomial model of "real option" value that can be used to value the mortgagee's put option. Although the specific application presented in Chapter 16 is to a "call option" instead of a "put option" (the option to buy, rather than to sell, the underlying asset), the same type of model can be applied to a put option.

\$100 million mortgage. However, in reality, the lender can only take possession of the property by going through the legal process of foreclosure, which is both time-consuming and costly. This gives the nonrecourse borrower additional potential leverage over the lender and suggests the possibility of the borrower engaging in **strategic default** behavior. The term **strategic default** refers to the borrower deliberately defaulting on the loan for the purpose of investment value maximization, even though they have the financial ability to make payments.

How can such behavior be in the borrower's interest? The presence of the deadweight burden of foreclosure costs can make it so. Consider a simple numerical example involving a nonrecourse mortgage. There is only one cash flow left on the loan, due 1 year from now, in the amount of \$10,600,000. The value of the underlying property at that time is currently uncertain, but has the following three possibilities:

- Scenario I: \$13 million
- Scenario II: \$11 million
- Scenario III: \$9 million

If foreclosure occurs, the costs of foreclosure will amount to \$2 million of deadweight burden (legal and administrative expenses paid to third parties). As the loan is nonrecourse, the lender would have to bear all of these costs in the event of foreclosure, although they could be recouped from the property sale proceeds.

Considering only the previous facts, it is easy to see what a wealth-maximizing borrower would do in each scenario. In scenario I, the borrower would not default. To default on the loan would be irrational because the lender could take the \$2 million costs of foreclosure plus the debt balance from the foreclosure sale proceeds, leaving the borrower with a net of only \$400,000 ($\$13,000,000 - \$10,600,000 - \$2,000,000$). On the other hand, by not defaulting, the borrower nets the full \$2.4 million difference between the property value and the loan balance. Thus, the NPV consequence of a decision to default would be $-\$2$ million. In effect, the borrower would needlessly pay the costs of foreclosure if they defaulted.

Now, skip to scenario III, the case in which the property is worth only \$9 million, which is less than the outstanding loan balance. (Such a circumstance is referred to as the loan being "under water" or an **underwater loan**.) In this case, it would be rational for the borrower to default, exercising his put option as described previously. If the borrower does not default, then he must pay the full \$10.6 million loan balance. As the property is only worth \$9 million, the decision not to default would have a negative NPV consequence of \$1.6 million. (Alternatively, you can view the decision to default as having a positive NPV impact of \$1.6 million relative to its only alternative.)

This would be an example of classical **rational default**, when defaulting on the loan is a positive NPV action that maximizes the NPV to the borrower across all the mutually exclusive alternatives available to him and hence satisfies the basic micro-level investment decision rule presented in Chapter 4.²⁷ Note that under scenario III, default would be rational even if there were no deadweight costs of foreclosure. The borrower might simply give the lender a deed in lieu of foreclosure in these circumstances. Otherwise, the \$2 million foreclosure costs would be absorbed by the lender, paid out of the property sale proceeds as the loan is nonrecourse, and the lender would net only \$7 million of the \$10.6 million owed.

Finally, consider the intermediate case, scenario II, in which the property is worth \$11 million. In this case, the loan is not under water: the property is worth more than the outstanding loan balance. However, it is not worth more than the sum of the outstanding balance plus the foreclosure costs.

²⁷ In general, it is not necessarily rational for the borrower to default as soon as the loan is under water if the loan is not yet mature. However, in the present simple example, the loan is due.

This is the situation in which strategic default becomes a profitable decision. The borrower can credibly threaten the lender with default, given the nonrecourse nature of the loan, and probably force the lender into a workout deal in which the debt is restructured to the borrower's advantage. Suppose, for example, that the borrower defaults on the loan and offers the lender the following workout deal. If the lender refrains from foreclosing on the property, the borrower will pay the lender \$9.1 million, instead of the \$10.6 million owed. Why would the lender take such a deal? Because by avoiding the \$2 million deadweight burden of the foreclosure costs, the lender will be better off by \$100,000. The borrower, of course, ends up \$1.5 million better off than if he had not defaulted. In effect, the two parties have split the \$1.6 million difference between the loan balance and the net value of the property in foreclosure, with the borrower getting the bulk of this difference because of his ability to force the lender to absorb up to \$1.6 million worth of foreclosure costs if he defaults.

The action described for the borrower in scenario II would be an example of a strategic default. In this case, the borrower did not have to default on the loan. The property was generating sufficient cash and value to service the debt. But by defaulting (or credibly threatening default), the borrower was able to improve her position by forcing the lender into a workout.

Of course, the real world is more complicated than the simple example described here. For one thing, both sides have **reputation effects** to consider. If the lender accepts the borrower's offer, the word may get out, and other borrowers with other loans may be tempted to play the same game with the lender. If instead, the lender calls the borrower's bluff and forecloses, he may develop a reputation as a "tough guy" who will not play such games. After all, if the lender forecloses in response to the borrower's default, the lender may lose up to \$1.6 million compared to what was owed under the loan, but the borrower also would lose \$400,000 compared to the case in which the borrower did not default. The borrower would also suffer the negative reputation effects of foreclosure, which could harm her future ability to borrow funds.²⁸ So the borrower may indeed be bluffing. In any case, the borrower would probably agree to pay quite a bit more than \$9.1 million to the lender to avoid formal foreclosure.

In practice, borrowers and lenders address workout situations in a variety of ways in the commercial mortgage business. In part, this may depend on the size of the loan and whether it is being held in the portfolio of the original lender. A portfolio lender may have more flexibility than the "special servicer" entity that governs the foreclosure decision in a CMBS pool of conduit loans.²⁹ While problem mortgages raise the specter of the deadweight costs of foreclosure, once a loan or borrower is in trouble, both opportunities and further dangers are presented from that point, with the resulting scope for creativity and entrepreneurship. Even as a mere possibility, actions such as strategic default need to be considered by lenders and mortgage investors in a probabilistic sense, for they can affect investment prospects and values.

19.6 CHAPTER SUMMARY

This chapter provides an overview of the U.S. mortgage industry, with a focus on commercial mortgages while also covering general legal and financial structures applicable to both commercial and residential mortgages in the U.S. We cover key concepts like foreclosure priority, typical mortgage covenants, and borrower defaults. The chapter highlights nonlitigious actions, such as workouts, as alternatives to foreclosure, and discusses the borrower's put option in nonrecourse loans, which can lead to strategic defaults. The impact of the COVID-19 pandemic on mortgage forbearance is also addressed, offering a timely analysis for understanding the complexities of mortgages from an investment perspective.

²⁸ Ciochetti and Vandell (1999) estimated that typical default transaction costs in permanent commercial mortgages average 5 percent and 15 percent of the outstanding loan balance for the borrower and lender, respectively.

²⁹ See Chapter 20 for discussion of CMBS conduit loans and the role of the special servicer.

KEY TERMS

- Whole loans
- Residential mortgage
- Commercial mortgage
- Secondary mortgage market
- Primary mortgage market
- Origination (issuance) of loans
- Securitization
- Commercial mortgage-backed
- Securities (CMBS)
- Permanent loans
- Construction loans
- Interest rate risk
- Floating interest rates
- Adjustable rate mortgages (ARMs)
- Depository institutions
- Mortgage bankers
- Mortgage brokers
- Take-out loan
- Collateralized debt obligation (CDO)
- Secured debt
- Mortgagee
- Mortgagor
- Promissory note
- Mortgage deed
- Deed of trust (security deed)
- Lien theory
- Title theory
- Power of sale
- Judicial sale
- Foreclosure
- Senior and junior debt
- Covenants
- Order of application of payments
- Acceleration clause
- Due-on-sale clause
- Loan assumption
- Wraparound loan
- Prepayment clause
- Defeasance
- Subordination clause
- Nonrecourse loans
- Default
- Litigious and nonlitigious actions
- Forbearance
- Deficiency judgment
- Bankruptcy
- Cramdown
- Chapter 11 bankruptcy
- Short sale

- Deadweight burden
- Deed in lieu of foreclosure
- Restructuring loans
- Workout of loans
- Limited liability
- Borrower's put option
- Credit risk (credit loss)
- Strategic default
- Underwater loans
- Rational default
- Reputation effects

STUDY QUESTIONS

Conceptual Questions

- 19.1 What determines whether a mortgage is classified as residential or commercial in the United States?
- 19.2 What are the major differences between the residential and commercial mortgage business in the United States?
- 19.3 Why is it that large commercial mortgage borrowers can often negotiate a customized loan from the issuer, while we small-fries trying to buy a house have to take a standard model?
- 19.4 When Brent takes out a mortgage to buy a \$10 million apartment building, is that transaction occurring in the primary or secondary mortgage market? Is Brent's position short or long in the debt asset?
- 19.5 Why are depository institutions predominant in the issuance of construction loans, while life insurance companies and pension funds are a larger presence in the permanent loan market? (Hint: Recall the maturity matching principle.)
- 19.6 What are the two parts (or legal documents) required in any mortgage?
- 19.7 What is a nonrecourse mortgage loan? Why do you think most permanent commercial mortgage loans are nonrecourse but residential mortgages, as well as construction loans, tend to be recourse?
- 19.8 Mortgage rates comparison
 - a. Which states would you expect to have slightly lower mortgage interest rates: title theory states or lien theory states? Why?
 - b. Answer the same question regarding power of sale versus judicial sale.
- 19.9 Bob is the mortgagee in a mortgage recorded on 1/31/2016 on the property at 1000 North Main Street in the amount of \$7,500,000. Sue has a mortgage on the property for \$1,000,000, recorded on 2/28/2018. Piet has a third mortgage for \$500,000, recorded on 3/31/2020. The property sells in foreclosure for \$9,000,000 in 2024, and there are \$1,000,000 in costs to third parties.
 - a. How much does each lien holder receive?
 - b. Which of the three lien holders is most likely to bring the foreclosure suit, and why?
 - c. How do your answers to (a) and (b) change if Sue's mortgage has a subordination clause in it and Piet's does not?
- 19.10 Why might a first mortgage holder not have the first priority on the proceeds from a foreclosure sale?
- 19.11 What is meant by the expression, "Redeem up, foreclose down"?

19.12 Order of application

- What is the normal order of application of payments in a mortgage?
- How does this order affect the remaining balance after an annual \$1 million payment is made (on time) on a mortgage with 10 percent simple annual interest and an outstanding balance of \$8 million just prior to the payment (i.e., what is the new balance after this payment is made)?
- Suppose the payment was late by one day, and the loan provisions call for a late-payment penalty fee of \$10,000.

19.13 Bob and Piet are 50/50 joint venture partners in a property investment in which they have borrowed \$10 million and both signed the mortgage. Apart from the real estate investment, Bob has a net worth of \$10 million, while Piet's net worth is 50 cents. They default on the loan when the property is worth \$8 million. Assuming foreclosure expenses of \$1 million, what are Bob's and Piet's respective net worths after the foreclosure, assuming the loan had a joint-and-several-liability clause but no exculpatory clause?

19.14 What is meant by acceleration of a loan? Name two common applications of acceleration clauses in a mortgage agreement.

19.15 When (and why) is it most valuable to the borrower for a loan *not* to have a due-on-sale clause in it?

19.16 Bob is the mortgagor on a \$5 million, 6 percent, interest-only first mortgage on the property at 1000 North Main. The loan has another five years to run before it comes due, and it has no due-on-sale clause. Since Bob bought the property, it has increased in value from \$7 million to \$10 million, while interest rates have increased from the 6 percent that had prevailed to 8 percent today. Sue now wants to buy the property from Bob for \$10 million, but she has only \$2 million available for a down payment and can only afford 7 percent interest on an \$8 million loan. Describe how Bob can clinch this deal by extending Sue a wraparound loan as a second mortgage.

19.17 Other things being equal, would you expect a lender to demand a higher interest rate on a loan with an exculpatory clause in it? Why?

19.18 What is cramdown? When does it occur?

19.19 Why might a lender employ forbearance, or avoid instituting a formal foreclosure procedure with a borrower in default?

19.20 Describe the nature and magnitude of the typical deadweight costs of foreclosure for commercial properties.

19.21 What are some of the specific options available to borrowers and lenders in a work-out process? (Hint: Describe various ways in which loan terms can be altered and/or additional equity capital injected into the property.)

19.22 Why are deficiency judgments often of relatively little value to the lender, even though the borrower has no exculpatory provision?

19.23 Why might a borrower give the lender a deed in lieu of foreclosure?

19.24 What is the relationship between the borrower's limited liability and the value of the borrower's put option?

19.25 What is strategic default, and how is it related to limited liability and the costs of foreclosure?

Quantitative Problems

19.26 Suppose \$10,600,000 is due and payable on a nonrecourse mortgage, and the property securing the loan is only worth \$9,000,000. What is the value to the borrower of her put option?

19.27 Consider a nonrecourse mortgage with one payment of \$10,600,000 due 1 year from now. The uncertain future is characterized by the following scenarios and probabilities:

Scenario I (70 percent probability): Property worth \$13,000,000
Scenario II (20 percent probability): Property worth \$11,000,000
Scenario III (10 percent probability): Property worth \$9,000,000

If foreclosure occurs, the costs paid to third parties will be \$2 million. U.S. government bonds maturing in 1 year are yielding 6 percent. If investors would require an expected return risk premium of 1 percent, what would this loan sell for today if scenario III would result in a deed-in-lieu and scenario II would result in a strategic default in which the difference between the borrower's and lender's extreme positions is split 50/50? What would be the loan's nominal yield, and what would be the present value cost of the credit risk?

Chapter Case: Keowee Ventures

In 2019, Keowee Ventures, owned by partners Kai and Lily, secured a commercial mortgage from Lakeside Bank to finance the expansion of their water recreational resort on the shores of Lake Keowee. The mortgage agreement included several covenants that affected their financial decisions in the following years.

In 2021, with interest rates at historically low levels, Kai and Lily wanted to refinance their loan. However, the mortgage agreement contained a specific clause that prevented them from paying off the loan early during the first three years after origination. By 2023, the financial landscape had shifted, and Keowee Ventures began experiencing financial difficulties. The Federal Reserve had raised interest rates, making their financial situation even more challenging. As they faced potential foreclosure, Keowee Ventures considered various options, including selling the resort, exploring defeasance, and catching up on missed payments.

Answer the Following Questions:

1. Which clause in the mortgage agreement prevented Keowee Ventures from refinancing their loan in 2021?
 - A) Due-on-Sale Clause
 - B) Acceleration Clause
 - C) Prepayment Clause
 - D) Good Repair Clause
2. What was the significance of the Due-on-Sale Clause in Keowee Ventures' mortgage agreement?
 - A) It allowed Keowee Ventures to sell the property without paying off the loan.
 - B) It required Lakeside Bank to lower the interest rate if the property was sold.
 - C) It allowed Lakeside Bank to demand full repayment of the loan if the property was sold.
 - D) It permitted Keowee Ventures to refinance the loan at any time.
3. Which clause could Lakeside Bank invoke when Keowee Ventures began missing payments, demanding that the entire loan balance be paid immediately?
 - A) Joint and Several Liability Clause
 - B) Acceleration Clause
 - C) Borrower's Right to Reinstate Clause
 - D) Lender in Possession Clause

4. When Keowee Ventures faced foreclosure, they consulted a professional about an alternative to directly paying off the mortgage. What is the process called where U.S. Treasury bonds are used as collateral instead of paying off the loan?
 - A) Defeasance
 - B) Subordination
 - C) Foreclosure
 - D) Loan Assumption
5. Which clause ensured that both Kai and Lily were fully responsible for the entire debt, even if one partner defaulted?
 - A) Good Repair Clause
 - B) Lender's Right to Notice Clause
 - C) Joint and Several Liability Clause
 - D) Future Advances Clause
6. If Keowee Ventures wanted to add new buildings to their resort in the future, which clause in the mortgage would allow Lakeside Bank to provide additional funding for this purpose?
 - A) Release Clause
 - B) Future Advances Clause
 - C) Acceleration Clause
 - D) Prepayment Clause
7. What does the Borrower's Right to Reinstate Clause allow Keowee Ventures to do?
 - A) Request a lower interest rate.
 - B) Stop loan acceleration by catching up on missed payments before foreclosure.
 - C) Transfer the mortgage to another borrower.
 - D) Pay off the loan early without penalties.

20 Mortgage Basics II

Payments, Yields, and Values

CHAPTER OUTLINE

- 20.1 Calculating Loan Payments and Balances
 - 20.1.1 Four Basic Rules
 - 20.1.2 Applying the Rules to Design Loans
- 20.2 Loan Yields and Mortgage Valuation
 - 20.2.1 Computing Mortgage Yields
 - 20.2.2 Why Points and Fees Exist
 - 20.2.3 Using Yields to Value Mortgages
- 20.3 Refinancing Decision
 - 20.3.1 Traditional Refinancing Calculation
 - *20.3.2 What Is Left Out of the Traditional Calculation
- 20.4 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- How to compute mortgage payments and balances for a variety of different types of loans, and how to creatively design your own customized loans.
- How to compute mortgage yields, and how to use mortgage yields to evaluate mortgages.
- The nature of the refinancing and prepayment decision, including the ability at some level to quantitatively evaluate this decision from a market value perspective.

The preceding chapter introduced you to the basic legal structure of mortgages and how they work. In this chapter, we will build on those concepts, as well as on the present value mathematical principles introduced in Chapter 2. Our objective here is to hone your understanding of how to compute loan payments, values, and yields, so that you can use this ability to creatively design and structure loans and effectively analyze debt investments. While the material in this chapter is rather technical and quantitative, it is also extremely practical and useful. Furthermore, it is much more fun than it used to be because much of the drudgery of formula memorization and number crunching is automated nowadays through the use of financial calculators and spreadsheets. This makes it easier for you to concentrate on the (more important) underlying economics.

20.1 CALCULATING LOAN PAYMENTS AND BALANCES

The first thing you need to master in handling a mortgage—whether as a borrower, lender, or investor—is the ability to calculate loan payments and balances based on the contracted loan terms. Now, you already have the tools to do this, based on what you learned in Chapters 2, 12, and 19. Chapters 2 and 12 provided the essential present value formulas and gave you some practice with the basic calculations. Chapter 19 covered the legal structure. Now, let's bring these together and explore how we can apply them more creatively.

20.1.1 FOUR BASIC RULES

(See Video 1 in Chapter 20 Supplements folder for a 9-minute tutorial on setting up loan payments & balance table in Excel.)

You may recall from Chapter 19 that one of the basic clauses present in all mortgages is the order of application of payments clause. This clause stipulates that payments received are applied to interest before principal. From this clause and the basic terms of the mortgage agreement, we derive the **four basic rules** for calculating loan payments and balances, which were already briefly mentioned in Chapter 12.

- Rule 1: The interest owed in each payment equals the applicable interest rate times the outstanding principal balance (i.e., the **outstanding loan balance**, or **OLB** for short) at the end of the previous period: $INT_t = (OLB_{t-1}) r_t$
- Rule 2: The principal amortized (paid down) in each payment equals the total payment (net of expenses and penalties) minus the interest owed: $AMORT_t = PMT_t - INT_t$
- Rule 3: The outstanding principal balance after each payment equals the previous outstanding principal balance minus the principal paid down in the payment: $OLB_t = OLB_{t-1} - AMORT_t$
- Rule 4: The initial outstanding principal balance equals the initial contract principal specified in the loan agreement: $OLB_0 = L$.

The following abbreviations are being used:

L = Initial contract principal amount (the loan amount)

r_t = Contract simple interest rate applicable for payment in period t

INT_t = Interest owed in period t

$AMORT_t$ = Principal paid down in the period t payment

OLB_t = Outstanding principal balance *after* the period t payment has been made

PMT_t = Amount of the loan payment in period t

The first of these rules derives from the basic nature of the mortgage agreement: the borrower agrees to compensate the lender for the use of the lender's money by paying interest. The second rule is the order of applications clause. (Under normal circumstances, we can ignore the expenses and penalties portion of the clause.) The third and fourth rules follow directly by definition. Together, these four rules define a complete system for determining all payments and the outstanding balance due on the loan at any point in time.¹

It's important to note that the term "applicable interest rate" in the first rule refers to the **contract interest rate** — the rate specified in the loan agreement—applied as a simple interest rate per payment period. In the United States, mortgage contract rates are generally presented as an annual interest rate. If the mortgage requires multiple payments in a year, you can calculate the simple per-period rate by dividing the annual rate by the number of payment periods per year. For example, a 12 percent mortgage with monthly payments actually applies a simple interest rate of 1 percent due at the end of each month.² Using the notation from Chapter 2: $r = i/m$, where i is the nominal annual rate and m is the number of payment periods per year. Also, keep in mind that the contract interest

¹ If the borrower pays more than required, the excess reduces the outstanding loan balance (OLB) faster, potentially resulting in an earlier payoff, though future payments may remain the same. These rules assume simple interest compounded on payment dates, as is typical in U.S. mortgages. However, interest can also be compounded more frequently than payments are made, with unpaid interest added to the principal balance—a process known as "accretion of principal." In such cases, the four rules still apply by aligning payment periods with the interest-compounding periods.

² Recall from Chapter 2 that this implies an effective annual rate (EAR) of $(1.01)^{12} - 1 = 12.68\%$, compounding the simple monthly rate at the monthly frequency. Also, we should note that in Canada residential and commercial mortgages use six-month compounding, not monthly, so first you compound to annual and then extract out the monthly rate for a monthly payment. For example, a 12 percent rate in Canada would be compounded like this $(1.06) \bullet (1.06) - 1 = 0.1236$ or 12.36% slightly lower than the U.S. result. For variable rate loans in Canada the compounding is monthly just as in the U.S.

rate in the mortgage does not necessarily equal the yield prevailing in the current secondary market, and it is most likely different from the ongoing interest rates in the primary mortgage market.

Similarly, the “initial principal” referred to in rule 4 is not necessarily equal to the net cash flow proceeds obtained by the borrower from the lender at the time the loan is made. For example, pre-paid interest (discount points) or origination fees may be deducted from the cash paid up front by the lender to the borrower.

20.1.2 APPLYING THE RULES TO DESIGN LOANS

The four rules described earlier practically cry out to be entered as formulas in the cells of a spreadsheet, from which you can use your creativity in designing a loan with payment and amortization patterns to your liking. Let’s examine a few of the more famous patterns.

Interest-Only Loan ($PMT_t = INT_t$, or equivalently: $OLB_t = L$, or $AMORT_t = 0$, for all t)

(See Video 2 in Chapter 20 Supplements folder for a 5-minute tutorial on setting up IO loan payments & balance table in Excel.)

Interest-only mortgages, also known as Zero-Amortization or Bullet mortgages, represent the oldest and most basic pattern of mortgage payments. In an interest-only mortgage, there is no amortization of principal; the outstanding loan balance remains constant throughout the life of the loan, with the entire original principal repaid in a single lump-sum “bullet” payment at the loan’s maturity date. As a result, the regular loan payments consist purely of interest. If the interest rate is fixed at a constant rate, then the regular payments will be fixed and level. This is the classical payment pattern of long-term corporate and government bonds, and it is not uncommon in commercial mortgages in the United States (with the difference that bond payments are traditionally semiannual whereas mortgage payments are monthly). Commercial construction mortgages are the most common type of interest-only loans in real estate.

The payment and interest profile of the interest-only loan is depicted graphically in Exhibit 20-1 for a \$1 million, 30-year fixed-rate mortgage at a 12 percent interest rate with monthly payments. Exhibit 20-2 shows the first and last scheduled payments. As 12 percent nominal annual interest equates to a simple monthly interest rate of 1 percent, the interest and payment amounts due each month on this loan are \$10,000. The last payment (at the end of month number 360) is \$1,010,000, including the last month’s interest and the entire original principal. The amounts in the cells of the loan schedule table are computed using the appropriate formulas from the previously described four basic rules, as indicated at the top of the table columns.³ It is important to note that with an adjustable-rate interest-only mortgage, the interest and payment amounts due each month will fluctuate according to the variable interest rate (Rule 1), while the principal balance and amortization will remain unaffected. This will be further illustrated in a later example under the topic of adjustable-rate mortgages.

The interest-only loan is very straightforward and easy to understand. It has the advantage to the borrower of regular payments that are less than those of an otherwise equivalent amortizing loan. Some borrowers may view as an advantage the fact that the entire regular payment is tax-deductible.⁴

³ Note that the four rules are applied as formulas in *three* of the four necessary columns in the spreadsheet (one column corresponding to each of the four variables: OLB, PMT, INT, and AMORT). If the rules were applied in formulas in *all four* columns, the system would be “circular.” Instead, one of the four columns is defined based on the defining characteristic of the loan type. For example, in Exhibit 20-1, we have chosen to define the PMT column based on the $PMT_t = INT_t$ (for all t) criterion of the interest-only mortgage. (Alternatively, we could have chosen to define any one of the other three columns based on the loan type definition, such as $AMORT_t = 0$ or $OLB_t = OLB_{t-1}$.)

⁴ Recall from Chapter 12 that only the interest portion of the debt service payment is tax deductible. However, recall from Chapter 13 that the tax-shelter value of debt to the borrower is often largely illusory, as much (if not all) of the taxes saved through the deduction are paid indirectly by the pass-through of the marginal lender’s taxes in the form of the difference between the market yield on taxed bonds and that on tax-exempt bonds.

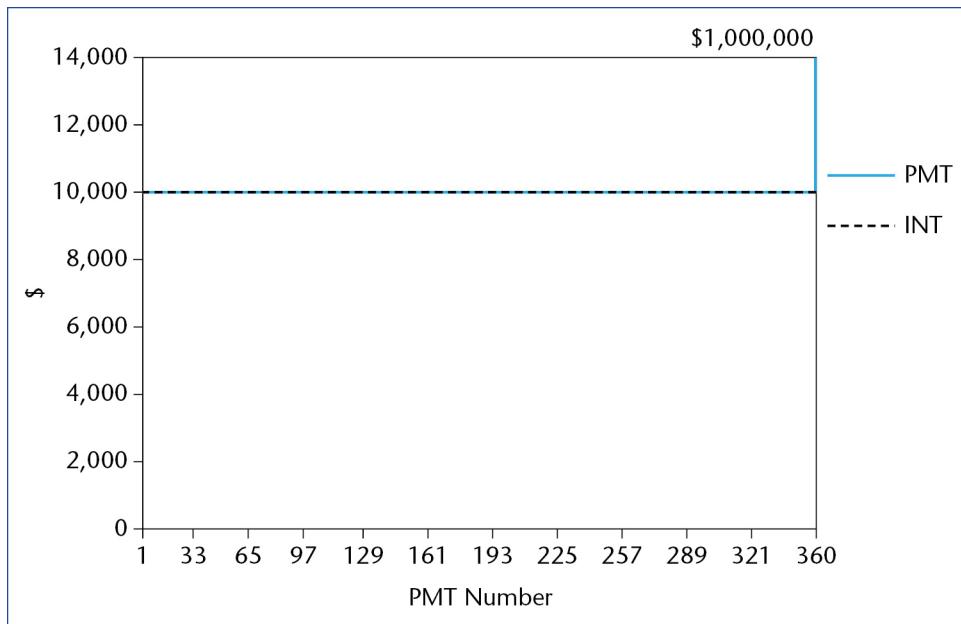


EXHIBIT 20-1 Interest-Only Mortgage Payments and Interest Component (\$1 million, 12 percent, 30-year fixed rate, monthly payments).

EXHIBIT 20-2

Interest-Only Mortgage Payments and Interest Component (\$1 million, 12 percent, 30-year fixed rate, monthly payments)

Month	Rules 3 & 4 OLB (Beg)	PMT	Rule 1 INT	Rule 2 AMORT	Rules 3 & 4 OLB (End)
0					\$1,000,000.00
1	\$1,000,000.00	\$10,000.00	\$10,000.00	\$0.00	\$1,000,000.00
2	\$1,000,000.00	\$10,000.00	\$10,000.00	\$0.00	\$1,000,000.00
3	\$1,000,000.00	\$10,000.00	\$10,000.00	\$0.00	\$1,000,000.00
...
358	\$1,000,000.00	\$10,000.00	\$10,000.00	\$0.00	\$1,000,000.00
359	\$1,000,000.00	\$10,000.00	\$10,000.00	\$0.00	\$1,000,000.00
360	\$1,000,000.00	\$1,010,000.00	\$10,000.00	\$1,000,000.00	\$0.00

On the other hand, because the principal is not paid down, an interest-only loan maximizes the total amount of interest paid over the life of the loan compared to a positively amortizing loan. Additionally, due to the repayment spike at the loan's maturity, the weighted average time until the loan payments are made is greater than that for an amortizing loan of the same maturity. This could have negative implications for some debt investors. For instance, if the loan has a fixed interest rate, its present market value becomes more sensitive to changes in market interest rates.⁵ The repayment spike at the end also forces the borrower to either refinance the loan or sell the property upon maturity, which

⁵ In general, the present values of cash flows that are more distant in the future are more sensitive to the average per-period discount rate than are the present values of more near-term cash flows. For example, the PV of \$110 1 year from now declines only from \$100 to about \$99 if the interest rate increases from 10 percent to 11 percent ($\$110/1.10$ versus $\$110/1.11$). However, the PV of a perpetuity of \$10 per year declines from \$100 to about \$91 with the same change in the interest rates ($\$10/0.10$ versus $\$10/0.11$).

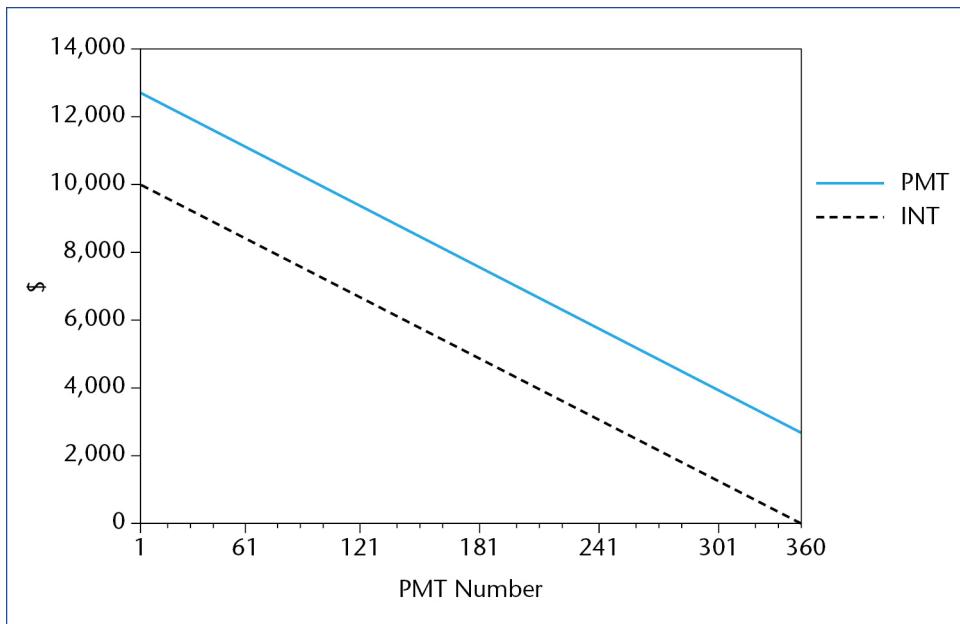


EXHIBIT 20-3 Constant-Amortization Mortgage (CAM) Payments and Interest Component (\$1 million, 12 percent, 30-year, monthly payments).

can be problematic if the property or debt market is unfavorable at that time. Indeed, this is the major drawback of the interest-only loan and the primary reason for the development of amortizing loans.

Constant-Amortization Mortgage (CAM) (AMORT_t = L/N, all t) The simplest way to solve the problem of the repayment spike at the end of the interest-only mortgage is to pay down a constant amount of principal in each loan payment. This results in a **constant-amortization mortgage**, or **CAM**. Such loans were used for a time in the 1930s when interest-only loans were causing havoc during the Great Depression (many a family farm was lost when the mortgage came due) and when persistent deflation resulted in declining rents and land values. As can be seen in Exhibit 20-3, the CAM is characterized by a declining payment pattern. As the loan balance is reduced by a constant amount each period, the interest owed falls by a constant amount as well given a fixed interest rate.

The payments on a CAM are computed by dividing the initial principal borrowed by the number of payments to compute the amortization amount per period, and then applying Rule 2 to compute the total payment due each period as the sum of the amortization and the interest computed based on rule 1 (Exhibit 20-4). Thus, for our \$1 million, 30-year example loan, the amortization each month is \$2,777.78, computed as \$1,000,000/360. Therefore, the first payment is \$12,777.78, considerably higher than the \$10,000 interest-only payment. After the first payment, the loan balance is reduced to \$997,222.22 by the application of rule 3, which in turn slightly reduces the amount of interest owed a month later in the second payment. This pattern continues until, after the 100th payment in our example, the CAM payment is less than that of the interest-only loan and continues to decline linearly to the end of the loan maturity. The outstanding loan balance also declines linearly.

In an economy free of persistent deflation, the declining payment pattern in the CAM can be undesirable for both borrowers and debt investors. From the borrower's perspective, this structure leads to excessively high initial loan payments, which may not align well with the income generation pattern of the property being financed. On the other hand, for mortgage investors, CAMs create the challenge of constant reinvestment of capital as the mortgage amortizes, potentially making reinvestment inconvenient or expensive. For these reasons, CAMs are not widely utilized in the United States today, though they may be suitable for certain unique circumstances.

EXHIBIT 20-4

Constant-Amortization Mortgage (CAM) Payments and Interest Component (\$1 million, 12 percent, 30-year, monthly payments)

Month	Rules 3 & 4 OLB (Beg)	Rule 2 PMT	Rule 1 INT	AMORT	Rules 3 & 4 OLB (End)
0					\$1,000,000.00
1	\$1,000,000.00	\$12,777.78	\$10,000.00	\$2,777.78	\$997,222.22
2	\$997,222.22	\$12,750.00	\$9,972.22	\$2,777.78	\$994,444.44
3	\$994,444.44	\$12,722.22	\$9,944.44	\$2,777.78	\$991,666.67
...
358	\$8,333.33	\$2,861.11	\$83.33	\$2,777.78	\$5,555.56
359	\$5,555.56	\$2,833.33	\$55.56	\$2,777.78	\$2,777.78
360	\$2,777.78	\$2,805.56	\$27.78	\$2,777.78	\$0.00

The Constant-Payment Mortgage (CPM) ($PMT_t = PMT$, a constant, for all t). This brings us to the classic, the constant-payment mortgage (CPM), which was briefly introduced in Section 12.4, in Chapter 12 in the printed core of this book. The CPM solves many of the problems associated with interest-only loans and the CAMs. Its consistent payments simplify budgeting for both borrowers and lenders, and it is more suitable for the typical rent growth pattern in income properties. The CPM also offers flexibility in balancing payment levels, amortization rates, and maturity periods. In fully amortizing loans, the CPM prevents the large repayment spike at maturity. When amortization is slower, payments can remain close to those of an interest-only loan, though this often requires long-term debt. A “hybrid” payment structure can be created by setting the loan maturity shorter than the amortization schedule. This results in a **balloon payment** at maturity, which is the large lump-sum payment made at the end of a loan to cover the remaining balance. The size of the balloon payment can be reduced by increasing the payment level (faster amortization). Although a CPM may not always have a fixed interest rate (the amortization can vary according to Rules 2 and 3), it is most common to combine the CPM payment pattern with a contractually fixed interest rate for the life of the loan.

The payments on the CPM are determined using the annuity formula introduced in Chapter 2, applied with the contractual interest rate, the initial contractual principal borrowed, and the number of payment periods specified in the contractual amortization of the mortgage: $PMT = L/[1 - 1/(1 + r)^N]/r$. Then rule 1 is applied to determine the interest components, and rule 2 is applied to back out the amortization component each period. (See Video 3 in Chapter 20 Supplements folder for a 4-minute tutorial on setting up IO loan payments & balance table in Excel.)

Continuing with our \$1 million, 12 percent, 30-year example in Exhibit 20-6, this results in a constant monthly payment of \$10,286.13 for the entire loan term. This is only slightly more than the \$10,000 payment for the interest-only loan, yet it fully amortizes the loan over 360 months without requiring a substantial balloon payment at maturity.

CPMs are characterized by an increasing amortization schedule. As demonstrated in Exhibit 20-6, after 15 years (180 payments), the remaining loan balance is still \$857,057.13, with less than 15 percent of the principal repaid during the first half of the loan’s term. This slow reduction in the loan balance early on is due to the fact that most of the payment is allocated to interest. In this example, more than half of the monthly payments are devoted to interest costs, with principal repayment becoming more substantial only after the 291st payment. From the 292nd payment onward, more than 24 years into the loan term, the portion of the payment applied to the principal finally exceeds that allocated to interest. In the later years of the CPM, principal repayment accelerates sharply, as indicated by the steep decline in the interest portion shown in the exhibit. It is also worth noting that

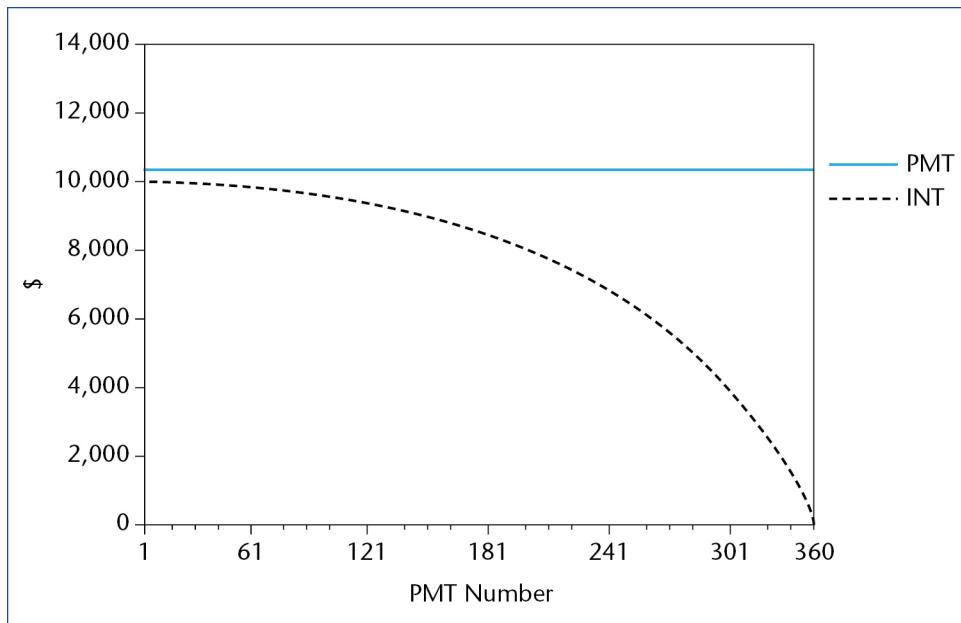


EXHIBIT 20-5 Constant-Payment Mortgage (CPM) Payments and Interest Component (\$1 million, 12 percent, 30-year, monthly payments).

EXHIBIT 20-6

Constant-Payment Mortgage (CPM) Payments and Interest Component (\$1 million, 12 percent, 30-year, monthly payments)

Month	Rules 3 & 4 OLB (Beg)	Rule 2 PMT	Rule 1 INT	AMORT	Rules 3 & 4 OLB (End)
0					\$1,000,000.00
1	\$1,000,000.00	\$10,286.13	\$10,000.00	\$286.13	\$999,713.87
2	\$999,713.87	\$10,286.13	\$9,997.14	\$288.99	\$999,424.89
3	\$999,424.89	\$10,286.13	\$9,994.25	\$291.88	\$999,133.01
...
180	\$858,755.70	\$10,286.13	\$8,587.56	\$1,698.57	\$857,057.13
...
291	\$516,039.66	\$10,286.13	\$5,160.40	\$5,125.73	\$510,913.93
292	\$510,913.93	\$10,286.13	\$5,109.14	\$5,176.99	\$505,736.94
...
358	\$30,251.34	\$10,286.13	\$302.51	\$9,983.61	\$20,267.73
359	\$20,267.73	\$10,286.13	\$202.68	\$10,083.45	\$10,184.28
360	\$10,184.28	\$10,286.13	\$101.84	\$10,184.28	\$0.00

as the interest component of the payments diminishes over time, the tax deductibility of the debt service is likewise reduced.

Adjustable Rate Mortgage (ARM) ($r_t \neq r_{t+s}$ for some s and t). Another way to improve the affordability of a mortgage for a borrower is to allow the contract interest rate in the loan to adjust periodically to changes in the interest rates prevailing in the debt market. This reduces the interest rate risk for the lender (or mortgage investor), making a lower interest rate possible.

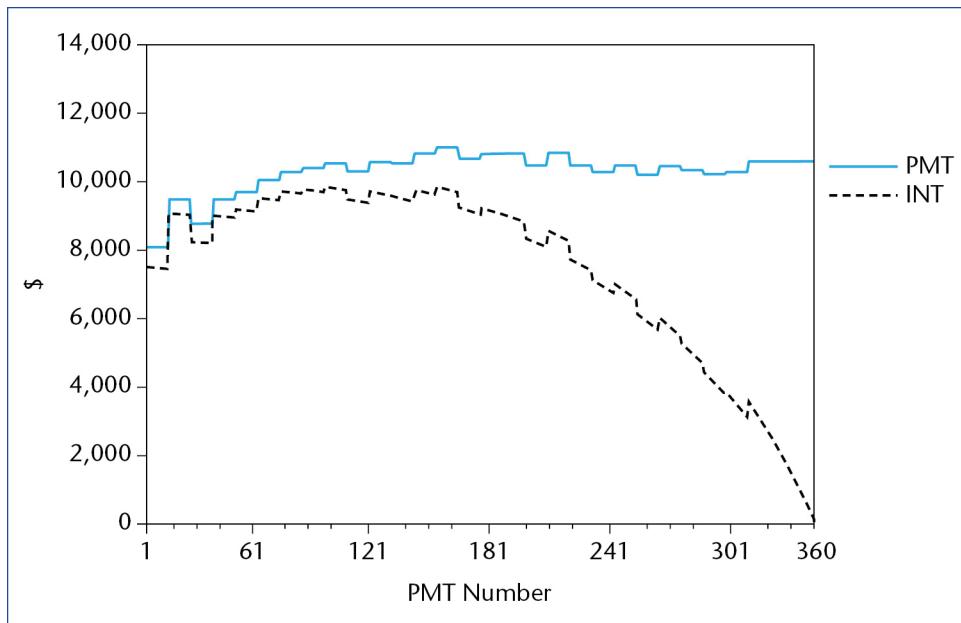


EXHIBIT 20-7 Adjustable Rate Mortgage (ARM) Payments and Interest Component (\$1 million, 9 percent initial interest, 30-year, monthly payments; one-year adjustment interval, possible hypothetical history).

The advantage of the ARM in reducing the initial interest rate is particularly strong for long-term mortgages during times when a steeply upward-sloping yield curve is prevailing in the bond market. **The yield curve** depicts the yield on bonds as a function of their maturity. With a steeply upward-sloping yield curve, short-term bonds are priced with a much lower yield than long-term bonds. This tends to occur when inflation is expected to increase in the long term, or when short-term real interest rates are temporarily depressed due to stimulative government monetary policy and/or low current demand for short-term capital (as during a macroeconomic recession). Although the ARM may be a long-term mortgage, it is like a chain of short-term fixed-rate loans linked together, because the interest rate can be adjusted at relatively short intervals.⁶ The ARM can therefore be priced similarly to short-term debt, with an interest rate based on the short-term end of the yield curve.

Continuing our previous numerical example, Exhibit 20-7 depicts a hypothetical \$1 million, 30-year ARM that might be available at the same time as the 12 percent loan we examined in the previous examples. This ARM has an initial interest rate of 9 percent, and the applicable interest rate adjusts once every year (12 payments), based on the prevailing interest rates on U.S. Treasury bonds with a 1-year maturity, plus a constant **margin** to reflect the greater default risk in the mortgage.⁷ In this case, the 1-year government bond is serving in the role of what is called the **index** for the ARM because it is governing the applicable interest rate on the ARM.⁸

⁶ Short-term commercial loans, such as construction loans, often have a “floating” interest rate, which may adjust frequently. Long-term (“permanent”) loans typically adjust no more frequently than once per year and they are based on a spread above SOFR, the Secured Overnight Financing Rate, based on transactions in the Treasury repurchase market, or other indices like LIBOR in Europe or COFI, a cost of funds index defined by a set of comparable lenders.

⁷ ARMs also often have “caps” or “ceilings,” placing limits on how far the applicable interest rate in the loan can adjust in any one move, or over the lifetime of the loan.

⁸ The index on an ARM must be a publicly observable rate that is not subject to manipulation by the lender. The most widely used index for residential loans is the U.S. Treasury bond with maturity corresponding to the adjustment interval in the ARM. Other widely used indexes include for commercial loans include cost of funds indices (COFI) relevant to the lender, SOFR, and the London Interbank Borrowing Rate (LIBOR).

EXHIBIT 20-8

Adjustable Rate Mortgage (ARM) Payments and Interest Component (\$1 million, 9 percent initial interest, 30-year, monthly payments; 1-year adjustment interval, possible hypothetical history)

Month	Rule 3 & 4 OLB (Beg)	PMT	Rule 1 INT	Rule 2 AMORT	Rule 3 & 4 OLB (End)	Applied Rate
0					\$1,000,000	
1	\$1,000,000.00	\$8,046.23	\$7,500.00	\$546.23	\$999,454	0.0900
2	\$999,454	\$8,046.23	\$7,495.90	\$550.32	\$998,903	0.0900
3	\$998,903	\$8,046.23	\$7,491.78	\$554.45	\$998,349	0.0900
...
12	\$993,761	\$8,046.23	\$7,453.21	\$593.02	\$993,168	0.0900
13	\$993,168	\$9,493.49	\$9,095.76	\$397.73	\$992,770	0.1099
14	\$992,770	\$9,493.49	\$9,092.12	\$401.37	\$992,369	0.1099
...
24	\$988,587	\$9,493.49	\$9,053.81	\$439.68	\$988,147	0.1099
25	\$988,147	\$8,788.72	\$8,251.03	\$537.68	\$987,610	0.1002
26	\$987,610	\$8,788.72	\$8,246.54	\$542.17	\$987,068	0.1002
...
358	\$31,100	\$10,605.35	\$356.61	\$10,248.74	\$20,851	0.1376
359	\$20,851	\$10,605.35	\$239.09	\$10,366.26	\$10,485	0.1376
360	\$10,485	\$10,605.35	\$120.23	\$10,485.12	0	0.1376

Ex-ante, it is impossible to know for certain what the future payments on the ARM will be beyond the initial year because it is impossible to forecast with certainty what the prevailing market yields will be for the index on the loan. Of course, a forecasted payment schedule can be computed based on an *assumed* forecast of future bond rates. Ex-post, the actual payment and balance history of the ARM is likely to look jagged, as in Exhibit 20-7, which assumes a particular (hypothetical) history for the government bond rate.

The interest rates applicable over the life of an ARM are likely to vary both up and down over time; however, they may tend to be higher on average than the initial interest rate, for two reasons. First, if the yield curve is steeply upward-sloping when the loan is issued, market yields even on short-term debt (such as the loan index) will tend to rise in the years after the loan is issued.⁹ Second, the initial rate on the ARM may be what is called a **teaser rate**. A teaser rate is an initial rate less than the current value of the index-plus-margin for the loan at the time the loan is issued. If the initial rate is a teaser, then the applicable rate on the loan will rise even if market interest rates (as represented by the rate on the loan index) remain the same.¹⁰

Although the applicable contract interest rate may vary unpredictably in the ARM, the loan will certainly fully amortize over its 30-year lifetime because the four rules will always be applied to determine the payments. To see how this works in practice, consider the payment and balance history

⁹ An abnormally steep yield curve is usually an indication that the capital markets are predicting a rise in future short-term interest rates. On the other hand, if the yield curve is slightly upward-sloping, with long-term rates, say, 100 to 200 basis points above short-term rates, then it will usually be approximately correct to forecast future short-term interest rates at a constant level equal to the current short-term rates.

¹⁰ For example, suppose the current 1-year T-bond is yielding 8 percent and an ARM with a one-year adjustment interval has a 2 percent margin. If the initial interest rate on the ARM were 9 percent, this would be a teaser rate, 1 percent below the fully indexed rate. In this case, even if the market interest rate remains at 8 percent for the T-bond, the ARM rate will rise to 10 percent at the first adjustment time.

depicted in the Exhibit 20-8. The initial interest rate on the loan is 9 percent. This gives the loan an initial monthly payment of \$8,046.23, found by applying the level annuity formula, $PMT(r,N,PV)$, with $r = 9\%/12 = 0.75\%$, $N = 30 \times 12 = 360$, and $PV = L = \$1,000,000$, the initial principal. After 12 payments and the application of the four basic rules, this leaves a remaining balance of \$993,168.

Now, at the 1-year anniversary of the loan, suppose that the 1-year U.S. Treasury bond yield turns out to be 8.99 percent. If the ARM has a margin of 200 basis points, then the applicable interest rate on the ARM will adjust from its previous 9 percent rate to a new rate of 10.99 percent. This is computed as the observed index rate on the anniversary date (8.99%) plus the 2 percent margin.¹¹ The payment for the second year of the ARM is then computed as $PMT(r,N,PV)$ with $r = 10.99\%/12 = 0.9158\%$, $N = 29 \times 12 = 348$, and $PV = \$993,168$, the outstanding loan balance at the beginning of the second year. This gives a monthly payment during the second year of \$9,493.49. This type of procedure is applied at each anniversary date or adjustment time in the loan.¹²

ARMs typically have at least slightly lower overall average expected yield-to-maturity than otherwise similar **fixed-rate mortgages (FRMs)**. This is because the yield curve is typically at least slightly upward-sloping, even in the absence of inflation fears or expansionary monetary policy. Of course, this does not imply that the borrower is getting something for nothing, for the borrower absorbs the interest rate risk that the lender avoids.

In our numerical example, the initial interest rate and payment are considerably lower in our 9 percent ARM than in the traditional fixed rate CPM. (Recall that the fixed-rate loan was at 12 percent, with the CPM payment at \$10,286.13, versus an initial payment of \$8,046.23 for the ARM.) However, with this much of a difference between their initial interest rates, it is not unlikely that the applicable interest rate and payment on the ARM will rise to levels above that of the equivalent FRM at some point during the life of the loan. For example, Exhibit 20-8 reveals that in the last year of the life of our example ARM, the applicable interest rate turned out to be 13.76 percent, which resulted in monthly payments that year of \$10,605.44, considerably above those on the corresponding FRM.¹³

Your Customized Loan. The types of mortgages and payment patterns previously described include the most common ones prevailing in the United States during the 1990s, but they obviously do not exhaust the possibilities.¹⁴ With a computer spreadsheet and the four basic rules, you can easily become a creative “financial engineer” and develop your own type of loan and payment pattern. Such creativity may not do you much good in trying to get a mortgage on your house, but remember that the commercial mortgage business is less standardized and often allows more customization and negotiation. Particularly in dealing with unique properties or investment circumstances, such as developments, turnarounds, and workouts, it may be useful to be able to be creative in designing loan terms and financing packages.

¹¹ Sometimes the relevant rate is specified as the *average* yield observed on the index during some window of time, such as the month prior to the anniversary date of the loan. This reduces the borrower's risk somewhat, as it tends to mitigate the danger associated with the anniversary date of the loan happening to fall at the time of a brief spike in market interest rates.

¹² In principle, it is possible for the payments in the ARM to be fixed even though the applicable interest rate varies. In this case, changes in the interest rate cause changes in the amortization of the contractual principal remaining on the loan, resulting either in changes in the loan maturity or the size of the balloon payment due at maturity. Fixed-payment ARMs are rare, however, in part because of the danger of “negative amortization” noted previously.

¹³ The ARM may have a prepayment clause in it entitling the borrower to refinance. (U.S. government regulations require residential ARMs to have the prepayment option.) However, the ARM payments will tend to fall when prevailing market interest rates fall, which reduces the value to the borrower of refinancing at that time. When market interest rates fall, the ARM rate will tend to automatically adjust downward as well. Thus, the prepayment option is typically less valuable in an ARM than it is in a FRM. As this option is less valuable to the borrower, so it is less costly to the lender. This is another aspect of the lender's reduced interest rate risk, which enables the ARM interest rate to be lower than rates on otherwise identical FRMs.

¹⁴ In Canada, we will see semiannual compounding used for the effective interest rates.

20.2 LOAN YIELDS AND MORTGAGE VALUATION

Investors, as we know, are interested in returns. This is no less true for debt investors than it is for equity investors. In the debt market in general, and in the mortgage market in particular, investment returns are generally referred to as **yields**.¹⁵ Now that you know how to compute loan payments and balances, you can easily learn how to compute loan yields. From there, it is a short step to loan valuation.

20.2.1 COMPUTING MORTGAGE YIELDS

The yield of a loan generally refers to its internal rate of return (IRR). Most commonly, this IRR is computed over the full remaining potential life of the loan, as if the loan would not be prepaid and the investor would hold the loan to maturity. To be more precise, this should be referred to as the **yield-to-maturity (YTM)** of the loan, although in common parlance, unless it is specified otherwise, the simple word *yield* usually refers to the YTM.

As an example, consider a \$1 million, 30-year CPM with an 8 percent annual interest rate. Such a loan would have monthly payments of \$7,337.65. The IRR of this loan computed as a simple monthly rate is found by solving the following equation for r .

$$0 = -\$1,000,000 + \sum_{t=1}^{360} \frac{\$7,337.65}{(1+r)^t}$$

Of course, the answer is $r = 0.667\%$, which equates to the nominal annual rate of $i = rm = (0.667\%)(12) = 8.00\%$. We know this because of how we determined the payment amount of \$7,337.65, by using the annuity formula with a nominal annual rate of 8 percent.¹⁶ In this case, therefore, the YTM of the loan is identical to its contract interest rate.

The yield on a mortgage will not always equal its interest rate. Suppose the previous \$1 million loan had a 1 percent origination fee. This means that the lender charges the borrower 1 percent of the loan amount up front, just to grant the loan. Such charges may also be referred to as prepaid interest or discount points (or just points, for short). They are normally quoted in points (a point in this sense refers to a percentage point of the outstanding loan balance).¹⁷ With a one-point origination fee, the lender will actually disburse to the borrower only \$990,000, even though the contractual principal (L) and the initial outstanding loan balance (OLB_0) is \$1 million.¹⁸ In this case, the YTM on the loan is found by solving the following equation for r .

$$0 = -\$990,000 + \sum_{t=1}^{360} \frac{\$7,337.65}{(1+r)^t}$$

The answer is $r = 0.6755$ percent, or a nominal annual rate of 8.11 percent. The effect of the origination fee is to increase the mortgage YTM by 11 basis points (or 0.11 percent) over the stated contract interest rate in the loan.

¹⁵ Recall from Chapter 3 that terminology in the bond and mortgage market is a bit different from that in the stock and equity market. Yield in the equity market typically refers to the income component of the periodic total return, or the current income or dividend value as a fraction of the current asset value. “Yield” in the debt markets typically refers to a multiperiod measure of the total return.

¹⁶ In this book, we will round monthly payments to the nearest cent, and present or future lump-sum values to the nearest dollar. A more exact monthly payment computation for a precise 8.00 percent yield is \$7,337.645739. In practice, the final payment on the loan would be adjusted slightly to make up for the effect of rounding.

¹⁷ The term “*basis points*” refers to one hundredth of a percent. One hundred basis points equals one point or 1 percent.

¹⁸ As soon as the borrower signs on the dotted line, he legally owes \$1 million, even though he receives only \$990,000.

Another way in which the YTM of a loan can differ from its contract interest rate is through the effect of the mortgage market on the value of the loan. Suppose the originator of this mortgage got the commitment and terms to this loan locked in a month before the closing of the loan transaction. During that month yields dropped in the mortgage market, so that the loan originator has an offer to sell this mortgage as soon as it closes for a value of \$1,025,000. The buyer of the loan in the secondary market is, in effect, offering to pay more for the loan than its current contractual outstanding balance (\$1 million), or par value. In effect, the mortgage market has caused the market value of the loan to differ from its par value or contractual OLB. The loan originator can make an immediate profit of \$25,000, or \$35,000 including the one-point origination fee. The buyer of the mortgage is, in effect, offering to make an investment in this mortgage with a YTM of 7.74 percent, found by solving the following equation for r (and multiplying by 12).

$$0 = -\$1,025,000 + \sum_{i=1}^{360} \frac{\$7,337.65}{(1+r)^i}$$

Contractual Interest Rates versus YTMs. In general, the YTM of a loan will differ from the loan's contractual interest rate whenever the current actual cash flow associated with the acquisition of the loan differs from the current outstanding loan balance (or par value) of the loan. At the time of loan origination in the primary market, this will result from discounts taken from the loan disbursement. In the resale of the loan in the secondary market, the YTM will reflect the market value of the loan regardless of the par value or contractual interest rate on the loan.

APRs and Effective Interest. The YTM from the lender's perspective at the time of loan origination is often referred to as the **annual percentage rate**, or **APR**. This term is used in consumer finance and residential mortgages, where it has an official definition based on the Truth-in-Lending Act.¹⁹ In common parlance, the YTM at the time of loan issuance is also often referred to as the **effective interest rate** faced by the borrower. However, it is important to recognize that the borrower may face additional transaction costs associated with obtaining the loan that are not reflected in the APR or YTM. In particular, costs for items that are required by the lender but paid to third parties, such as appraisals and title insurance, are not included in the APR calculation. These costs may differ across lenders, so the actual lowest cost of capital for the borrower will not necessarily come from the lender with the lowest APR.

YTMs and Expected Returns. The YTM of a loan may differ from the investor's expected total return (and therefore also from the borrower's expected cost of capital) for two reasons. First, the YTM computed as before ignores the possibility of default on the loan by the borrower. The contractual cash flows specified in the loan agreement are assumed in computing the yield. More realistically, there is normally some ex-ante probability that default will occur at some point in the life of the loan. This causes the realistic ex-ante mean return faced by the investor to be less than the default-free YTM, which is the way loan yields are measured and quoted in the market.²⁰

Second (even ignoring the possibility of default), if the loan has a prepayment clause or a due-on-sale clause, then the borrower may choose to pay the loan off before its maturity. This is common in residential mortgages, and not uncommon in commercial mortgages. The yield (IRR) over the realistic expected life of the loan (until the borrower prepays) will differ from the YTM whenever there is a **prepayment penalty** and whenever the YTM differs from the contractual interest rate on the loan.

¹⁹ See Federal Reserve Board, Regulation Z, 12 C.F.R. 226 as amended. This regulation does not apply to most commercial mortgages. Note that the APR is generally rounded to the nearest eighth of a point.

²⁰ Default risk will be discussed in more depth in the next chapter. Keep in mind, however, the basic principle introduced in Chapter 19, that the value of the loan to its holder equals the value of an otherwise equivalent default-free loan (such as a government bond) minus the value of the borrower's put option represented by the borrower's limited liability. Thus, option valuation theory can be used to value default-risky debt, and the YTMs observed in the debt market reflect such valuation.

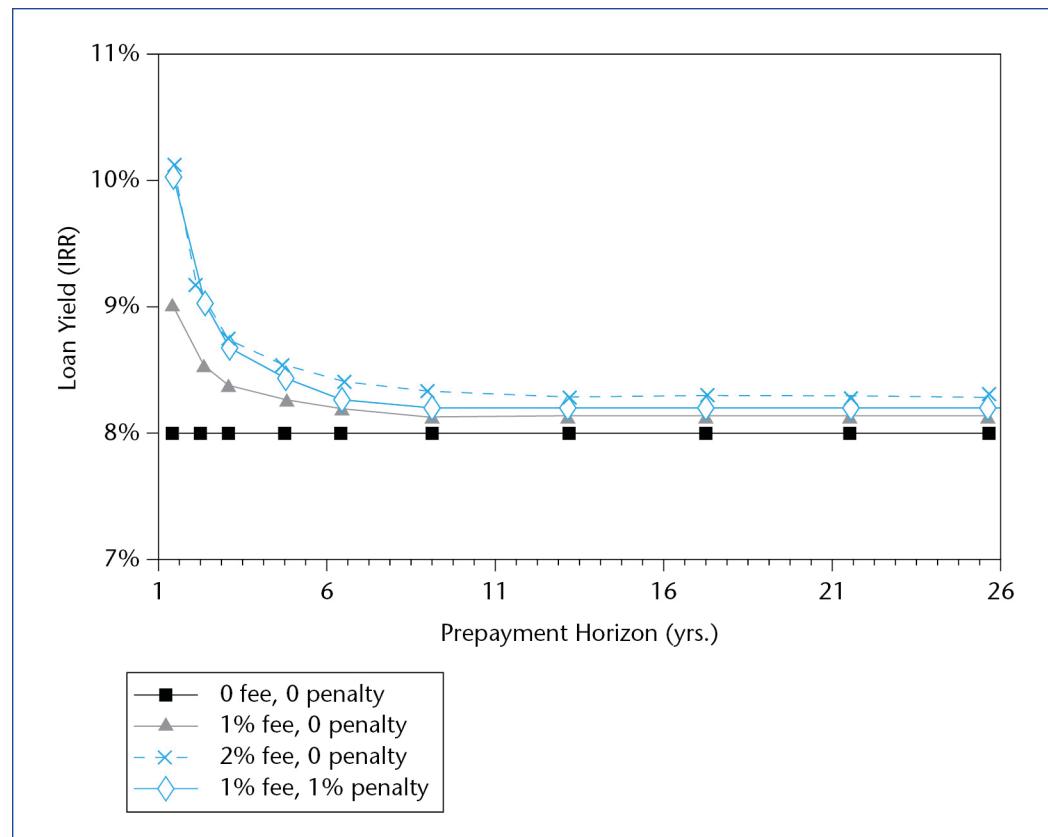


EXHIBIT 20-9 Effect of Prepayment on Loan Yield (8 percent, 30-year).

At the time of loan origination in the primary market, the YTM differs from the contractual interest rate only due to disbursement discounts, so the expected yield to the lender (or effective interest rate to the borrower) generally *increases* the shorter the time until the realistic prepayment horizon. This effect is magnified if there is also a prepayment penalty.

As an example, let's return to our previous \$1 million, 30-year, 8 percent loan with the one-point origination fee. We saw that its YTM at origination was 8.11 percent. Now, suppose the borrower is expected to prepay the loan after 10 years. The yield of the loan with an expected prepayment horizon of 10 years is 8.15 percent, computed by solving the following equation for r (and multiplying by 12).²¹

$$0 = \$990,000 + \sum_{t=1}^{120} \frac{\$7,337.65}{(1+r)^t} + \frac{\$877,247}{(1+r)^{120}}$$

²¹ Recall from Chapter 2 that a business calculator solves the following problem:

$$0 = -PV + \sum_{t=1}^N \frac{PMT}{[1 + (i/m)]^t} + \frac{FV}{[1 + (i/m)]^N}$$

Thus, to solve the problem on a business calculator, first enter the contractual loan terms ($N = 360$, $i = 8\%$, $PV = 1000000$, $FV = 0$, and compute $PMT = -7337.65$). Then change the number of payments to that in the expected horizon ($N = 120$) to compute the OLB as $FV = -877247$, before changing the amount in the PV register to reflect the fee: $PV = 990000$. It is important to change the amount in the PV register last, just before computing the yield using the i register. This is a general rule. For example, if the loan had also had a prepayment penalty, it would be important to compute the amount of the prepayment penalty with the original contract principal in the PV register, and include the prepayment penalty in the FV register amount, before changing the value in the PV register to reflect any disbursement discount.

EXHIBIT 20-10

Yield (IRR) on 8 percent, 30-year CP-FRM

Loan Terms	Prepayment Horizon (yrs.)						
	1	2	3	5	10	20	30
0% fee, 0 pen	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%
1% fee, 0 pen	9.05%	8.55%	8.38%	8.25%	8.15%	8.11%	8.11%
2% fee, 0 pen	10.12%	9.11%	8.77%	8.50%	8.31%	8.23%	8.21%
1% fee, 1% pen	10.01%	9.01%	8.67%	8.41%	8.21%	8.13%	8.11%

The resulting 8.15 percent yield is more realistic than the 8.11 percent YTM if this loan would in fact more likely be repaid in 10 years. The 8.15 percent yield is also the YTM of a 10-year mortgage with a 30-year amortization rate. Such a loan would have a mandatory balloon payment of \$877,247 at the end of its 10-year maturity.²²

The shorter the prepayment horizon, the greater the effect of the disbursement discount on the realistic yield of the mortgage. Prepayment penalties cause a similar (though slightly smaller) effect. The effect is hyperbolic, not linear, with the impact on the yield much greater the shorter the prepayment horizon. This is seen in Exhibit 20-9 and Exhibit 20-10, which show the effect on the yield of a 30-year, 8 percent mortgage caused by various prepayment horizons and loan terms. For example, with one point of disbursement discount the yield would be 8.25 percent over a 5-year horizon, or 8.55 percent over a 2-year horizon. If the loan had two points of disbursement discount instead of one, its yield over the 10-year horizon would be 8.31 percent instead of 8.15 percent. If it had one point of prepayment penalty instead of the additional point of disbursement discount, the 10-year yield would be 8.21 percent.²³

20.2.2 WHY POINTS AND FEES EXIST

In the commercial mortgage industry, origination fees are often charged to compensate mortgage brokers who find and filter loan applications for the financial institutions that supply the long-term capital. In general, loan originators face overhead and administrative costs that may be attributed to the loan origination process *per se*, and which therefore differ from the opportunity cost of the capital that is being invested in the loan. While the latter cost accrues over the life of the loan, the former cost is sunk at the time the loan is issued. Up-front fees are a way for the loan originator to make some profit while providing some disincentive against early prepayment of the loan by the borrower. (Recall the hyperbolic relation between the yield and the prepayment horizon when there is a disbursement discount.) Up-front fees can also be used as a trade-off against the level of the regular loan payment: greater origination fees and discount points allow lower regular loan payments for the same yield. This results in the familiar **mortgage menu** typically offered by residential lenders, in which discount points are traded off against the interest rate on the loan.

²² The OLB on the loan can be computed on a business calculator in either of two ways. You may enter the number of payments *already made* in the *N* register and then compute the future value (FV) of the loan (with the initial principal amount in the PV register). Alternatively, with the loan's original liquidating payment at maturity in the FV register (zero for a fully amortizing loan), it is equivalent to enter the number of payments *yet to be paid* in the N register, and then compute the present value (PV). That these two approaches are mathematically equivalent can be seen by manipulating the mortgage math calculator equation [equation (12) in Chapter 2]. Letting q be the number of previous payments and Tm be the original number of payments in the fully amortizing loan maturity, it is equivalent to solve this equation either for PV with $N = Tm - q$ and $FV = 0$, or for FV with $N = q$ and $PV = P$, the original contract principal.

²³ It is computationally useful to note that when the disbursement discount and prepayment penalty are quoted in points, that is, as a fraction of the outstanding loan balance or par value of the loan, then there is no need to know the dollar magnitude of the loan or its payments in order to compute its yield. You can substitute any convenient loan amount (such

EXHIBIT 20-11

At the Other End of the Loan's Lifetime, Prepayment Penalties are Rare in Residential Mortgages in the United States, But are More Widespread in Commercial Mortgages. They are Used to Discourage Prepayment, in Effect, by Reducing the Net Value of the Borrower's Prepayment Option. In the Extreme, Prepayment Penalties May be Set so High as to Effectively Eliminate the Value of this Option²⁴

Discount Points	Interest Rate	Monthly Payment
0	8.15%	\$7,444.86
1	8.00%	\$7,337.65
2	7.85%	\$7,230.58
3	7.69%	\$7,124.08

For example, suppose the originator of the previously described \$1 million mortgage with the 30-year amortization rate expects (or requires) that the loan will be paid back in 10 years (i.e., with a balloon payment and no penalty). The following combination of disbursement discount points and contract interest rate will all provide the same 8.15 percent 10-year yield, with the resulting indicated difference in the monthly payments (Exhibit 20-11).

20.2.3 USING YIELDS TO VALUE MORTGAGES

Now that you know how to compute mortgage yields, you can use this knowledge as a valuation tool. Investors buy mortgages (or issue them as loan originators) in order to earn returns, as with any other investment. We have seen how the mortgage yield is related to the return the mortgage investor can expect. As with any investment, the return (and therefore the yield) is an inverse function of the price of the asset (holding the future cash flows constant). For a given loan (i.e., a given set of contractual future cash flows), a given yield will therefore correspond to a certain price, and vice versa. In effect, yields become a convenient way to measure and quote the prices of mortgages prevailing in the debt market. By applying the relevant market yield to a given mortgage, you can ascertain the market value of that mortgage, and what it is likely to sell for in the secondary market. In general, for a given required yield, you can determine how much a specified mortgage is worth.

As an example, consider once again our \$1 million, 8 percent, monthly-payment, 30-year- amortization, 10-year-maturity balloon loan. How much is this loan worth at closing if the market yield at that time is 7.50 percent? The answer is \$1,033,509, computed as²⁵

$$\$1,033,509 = \sum_{t=1}^{120} \frac{\$7,337.65}{(1.00625)^t} + \frac{\$877,247}{(1.00625)^{120}}$$

This is just the inverse of the yield calculation problems we did in the last section. There, we knew the present cash flow value (the PV amount), and solved for the implied return. Now we know the

as \$1) to compute the yield for any loan so described. This is a consequence of the fact that the present value equation is homogeneous of degree one in the cash flows, as noted in Chapter 2.

²⁴ The prepayment penalties described in this chapter are defined as a simple, constant amount of “points” (basis points) throughout the life of the loan, that is, a dollar amount that is always a constant proportion of the current outstanding loan balance. However, in commercial mortgages, it is common for prepayment penalties to be defined based on “yield maintenance” provisions.

²⁵ Here, the market yield (discount rate) has been quoted as a **mortgage-equivalent yield (MEY)**, so the nominal annual rate of 7.50 percent corresponds to a simple monthly rate of $7.5\% / 12 = 0.625\%$. In practice, the mortgage market and bond market are tightly integrated, and market yields are usually quoted in **bond- or coupon-equivalent terms (BEY or**

required return, so we must solve for the implied present value. This is nothing more (or less) than a DCF valuation problem similar to those we introduced in Chapter 4, only now we are valuing a debt asset rather than an equity asset.²⁶ The discount rate is the yield observed in the debt market for loans with similar risk and payment timing patterns as the subject loan. The implied loan value is that at which the loan can be traded with $NPV = 0$ based on market value.²⁷

Similarly, if both the yield and the worth of the mortgage are specified, you can determine the required payments and hence the contract interest rate and/or discount points that are required for a given par value or contractual principal amount. For example, suppose the market yield on the day of closing of the 8 percent, \$1 million loan is 8.5 percent instead of 7.5 percent. How much must the loan originator charge the borrower in the form of discount points if the lender is not to do a negative NPV deal from her perspective, based on market value?

To answer this question, we solve the same type of loan valuation equation as before, to see that the loan will be worth \$967,888 on closing:²⁸

$$\$967,888 = \sum_{t=1}^{120} \frac{\$7,337.65}{(1.0070833)^t} + \frac{\$877,247}{(1.0070833)^{120}}$$

As this value is \$32,112 less than the initial principal, that is the amount by which the lender must discount the up-front disbursement. In other words, 3.2112 points of prepaid interest or origination fees must be charged.²⁹

20.3 REFINANCING DECISION

If a loan has a **prepayment option**, the borrower can choose to pay the loan off early to take advantage of favorable interest rate movements, **refinancing** the old loan with a new, cheaper loan. This refinancing decision can be evaluated by comparing two loans, the existing (“old”) loan and a “new” loan that would replace it. Traditionally, this comparison is made using the classical DCF methodology we have just been describing. In this section, we will first present this traditional approach, then we will explore something important that is left out of the traditional picture, namely, the pre-payment option value in the old loan.

20.3.1 TRADITIONAL REFINANCING CALCULATION

The traditional approach involves a simple DCF valuation of the two loans. Here we will label the DCF-based present values of the two loans $PV(CF^{OLD})$ and $PV(CF^{NEW})$, for the old and new loans, respectively. The $PV()$ function here refers to the DCF procedure, and CF^{OLD} represents the

CEY). Yields quoted in bond terms must be converted to MEY before application to monthly cash flows. In the previous example, if the market yield is 7.5 percent MEY, then the observed market yield in the bond market (for a bond of equivalent risk and maturity) must be 7.62% BEY, as $[1 + (0.0762/2)]^2 = [1 + (0.075/12)]^{12}$. See Chapter 2 for a discussion of mortgage-equivalent and bond-equivalent yields and effective annual rates (EARs).

²⁶ As noted, one difference here is that the contractual cash flows are used in the numerators on the right-hand side of the valuation equation. In Chapter 4, the realistic *expectations* of the cash flows in each period were used in the numerators. Considering the probability of default, the contractual cash flow amounts are generally higher than the realistic expectations of these cash flows. Thus, the discount rate (the market yield) in the corresponding denominators is gross of the expected credit loss to the lender in the event of default.

²⁷ With reference to Chapter 14, note that the valuation described here computes the market value (MV) of the loan, which equates to the investment value (IV) only for marginal investors.

²⁸ $8.5\%/12 = 0.070833\%$

²⁹ This shows how leaving the final determination of the loan origination fee flexible until the closing can enable the latest changes in the secondary mortgage market to be reflected in the loan terms.

remaining cash flows in the old loan. CF^{NEW} represents the future cash flow stream in the new loan (after its initial disbursement to the borrower). In order to make an apples-to-apples comparison, both loans must be evaluated over the same time horizon (including the same remaining term on the loan), and for the same loan amount.³⁰ Furthermore, both loans must be evaluated using the current opportunity cost of capital as the discount rate. With the two loans evaluated in this manner, the NPV of refinancing for the borrower is defined as the value of the old liability minus the value of the new liability, quantified as $PV(CF^{OLD}) - PV(CF^{NEW})$, less any transaction costs the borrower faces in the refinancing deal.

This procedure is equivalent to calculating the net incremental difference in cash flows each period under the new loan compared to the old, and summing the present values of these periodic savings to the borrower. Using the same discount rate (the current opportunity cost of capital), the difference in the present values of the cash flows equals the present value of the differences in the cash flows:³¹

$$PV(CF^{OLD}) - PV(CF^{NEW}) = PV(CF^{OLD} - CF^{NEW})$$

Normally, the savings from refinancing will occur in the regular monthly payments, while the last cash flow in the analysis time horizon may involve an incremental cost to the borrower, as the new loan may at that time have a larger outstanding balance than the old loan would have had.

To implement this DCF procedure in practice, we need to define specifically how to determine the common time horizon applied to both loans, the amount of the new loan, and the discount rate. The common time horizon should be the expected time until the old loan would likely be paid off in the absence of refinancing, which at the latest is the maturity date of the old loan.³² This may be earlier than the maturity of the replacement loan, in which case the replacement loan must be evaluated with expected cash flows corresponding to such early prepayment of the replacement loan. In order to keep the amount of debt constant, the new loan should be evaluated for an amount such that the actual cash disbursement the borrower receives is just sufficient to pay off the old loan exactly. Thus, if the new loan has a disbursement discount, then the contractual principal borrowed on the new loan must exceed the current OLB plus prepayment penalty on the old loan. Finally, the opportunity cost of capital used as the discount rate in the analysis should be determined as the yield on the new loan, computed over the common time horizon. If the new loan has any disbursement discounts, or if the new loan has a prepayment penalty and the common time horizon is prior to the new loan's maturity, then this discount rate will, of course, exceed the contract interest rate in the new loan.

Shortcut Procedure The preceding conditions on the discount rate and the amount of the new loan imply that a shortcut exists to quantify the difference between the new and old loan values. In effect, *we do not need to quantify the amount of the new loan or its payments*. Defining the opportunity cost of capital as the yield on the new loan assures that $PV(CF^{new})$ equals the cash disbursement to the borrower on that loan. And the condition on the new loan amount requires that this cash disbursement must equal the amount required to pay off the old loan, an amount we will label

³⁰ This condition is necessary in order to avoid mixing two distinct decisions: the refinancing decision and a decision to change the degree of leverage on the equity investment. The importance of keeping these two decisions separate can be seen by recourse to an extreme example. Suppose a \$1 million loan is to be refinanced with a \$500,000 loan. The borrower will be getting rid of an old liability that has twice the magnitude of the new liability. Simply comparing the values of the two liabilities, it would appear that the borrower would be approximately \$500,000 better off no matter what the difference in interest rates! Obviously, this would be misleading from the perspective of evaluating the refinancing decision per se. To isolate the NPV of this decision alone, we hold constant for analysis purposes the amount of debt in comparing the old loan with its replacement loan. Of course, in practice, the investor may choose to change the amount of debt at the same time as the refinancing.

³¹ This identity follows directly from the distributive law of multiplication and addition: $ab + ac = a(b + c)$.

³² It may be earlier than the maturity if the loan has a due-on-sale clause.

OLB^{OLD} . Thus, the conditions described above imply $PV(CF^{new}) = OLB^{old}$. The refinancing NPV can thus be redefined (apart from transaction costs) as: $PV(CF^{old}) - OLB^{OLD}$.

This result has an important conceptual implication. Fundamentally, the refinancing decision is not a comparison of two loans. Rather, it is a decision simply regarding the old loan: *Does it make sense to exercise the old loan's prepayment option?* It does not matter whether the old loan would be paid off with capital obtained from a new loan, or additional equity, or some combination of debt and equity. Thus, the refinancing decision is simply a comparison of the current liability value of the old loan with the cash that would currently be required to pay off the old loan. A new loan is necessary in the analysis only as a (possibly hypothetical) source for ascertaining the current relevant opportunity cost of capital.

An Example NPV Calculation. The best way to understand the traditional refinancing valuation is to consider a specific numerical example. Let's go back to our previously described 8 percent, \$1 million mortgage with 30-year amortization and 10-year maturity at origination. Suppose this loan was issued four years ago and has a prepayment option with a penalty of two points. Suppose that if not refinanced, this loan would probably be held to its maturity, six more years. Now, suppose that new loans are available today with a maturity of six years at an interest rate of 7 percent, with one point of disbursement discount up front, and an amortization rate of 30 years. What is the net value of refinancing the old loan at this point in time (apart from transaction costs)?

Step 1: Old Loan Liquidating Payment Let's begin by calculating what it would take to pay off the old loan today. The outstanding balance on the old loan after four years (48 payments) is \$962,190, and when we add two points of prepayment penalty, this gives a liquidating payment of $1.02 \times \$962,190 = \$981,434$.³³ This is the amount we labeled OLB^{OLD} .

Step 2: Opportunity Cost of Capital let's compute the relevant cost of capital as the yield on the new loan over the remaining maturity of the old loan. We don't need to know the loan amount to do this. The loan would have monthly payments based on a 360-month level annuity at a simple interest rate of $7\% / 12 = 0.5833\%$ per month. For every dollar of loan amount, this is a monthly payment of 0.006653 dollars: $PMT(0.07/12, 30 * 12, 1) = 0.006653$. The balloon at the end of the six-year maturity on the new loan would be \$0.926916 per dollar of loan amount:

$$\begin{aligned} PV & \left(0.07/12, 24 \times 12, \left[PMT(0.07/12, 30 \times 12, 1) \right] \right) \\ & = FV \left(0.07/12, 6 \times 12, \left[PMT(0.07/12, 30 \times 12, 1) \right] \right) = 0.926916 \end{aligned}$$

Considering the one-point disbursement discount up front, this gives the new loan a yield over the 6-year horizon of 7.21 percent, as³⁴

$$0.99 = \sum_{t=1}^{72} \frac{\$0.006653}{\left[1 + (0.0721/12) \right]^t} + \frac{\$0.926916}{\left[1 + (0.0721/12) \right]^{72}}$$

This would be the yield (or effective interest rate) no matter what the amount of the new loan. Thus, 7.21 percent is the current opportunity cost of capital (OCC) relevant to our refinancing calculation.

Step 3: Present Value of the Old Loan Liability let's compute the present value of the old loan's remaining cash flows using the 7.21 percent OCC we just computed in step 2 as the

³³ $\$962,190 = PV(0.08/12, 26 \cdot 12, 7337.65) = FV(0.08/12, 48 \cdot 12, 7337.65)$, where $\$7,337.65 = PMT(0.08/12, 30 \cdot 12, 1, 000,000)$

³⁴ $0.0721 = 12 \cdot RATE(6 \cdot 12, 0.006653, 0.99, 0.926916)$

discount rate. The old loan has regular monthly payments of \$7,337.65 and a balloon payment of \$877,247 at its maturity, six years from now (10 years from the issuance of the loan).³⁵ Thus, the present value of this liability to the borrower is now \$997,654, computed as³⁶

$$97,654 = \sum_{t=1}^{72} \frac{\$7,337.65}{[1 + (0.0721/12)]^t} + \frac{\$877,247}{[1 + (0.0721/12)]^{72}}$$

This is the value we labeled $PV(CF^{OLD})$. Traditionally, this amount is viewed as the present value of the benefit of the refinancing to the borrower, as it is taken to represent the current value of the liability that would be removed by paying off the old loan.

Step 4: Computing the NPV we can compute the traditional NPV of refinancing to the borrower before considering transaction costs. This is simply the present value of the benefit computed in step 3 less the amount of the liquidating payment computed in step 1:

$$NPV = \$997,654 - \$981,434 = \$16,220$$

This is the value of $PV(CF^{OLD}) - LB^{OLD}$, including the prepayment penalty in LB^{OLD} .

The Long Route: Specifying the New Loan Let's confirm that this shortcut gives us the same NPV as if we specified the replacement loan amount. From step 1, we know we need \$981,434 of cash disbursement from the new loan in order to pay off the old loan. As the new loan has a one-point disbursement discount, we would have to borrow \$991,348 as the initial contractual principal in the new loan, computed as \$981,434/0.99. At 7 percent interest, this gives the new loan a monthly payment of \$6,595.46 and a balloon of \$918,896 after six years.³⁷ We already know from step 2 that the yield on this new loan, over the 6-year horizon, is 7.21 percent, including the effect of the one-point disbursement discount. Now, what is the present value of this new loan discounted at the 7.21 percent rate? Of course, it is \$981,434! The yield, by definition, is the discount rate that causes the loan's future cash flows to have a present value exactly equal to its time-zero cash disbursement.³⁸

Thus, we can immediately confirm that the difference in the two loan values is \$997,654 – \$981,434 = \$16,220. In other words, $PV(CF^{OLD}) - PV(CF^{NEW}) = PV(CF^{OLD}) - OLB^{OLD}$. We don't need to deal with the new loan, except as a means to ascertain the relevant current cost of capital.³⁹

*20.3.2 WHAT IS LEFT OUT OF THE TRADITIONAL CALCULATION

Prepayment Option Value

According to the traditional analysis performed earlier, our example loan should be prepaid, as long as the transaction costs involved in obtaining the necessary capital are less than \$16,220. Suppose, for example, that such transaction costs would be \$10,000 (approximately 1 percent

³⁵ $\$877,247 = PV(0.08/12, 20 \cdot 12, 7337.65) = FV(0.08/12, 10 \cdot 12, 7337.65)$

³⁶ $\$997,654 = PV(0.0721/12, 6 \cdot 12, 7337.65, 877247)$

³⁷ $\$991,348 = PV(0.07/12, 6 \cdot 12, 6595.46, 918896)$

³⁸ $\$981,834 = PV(0.0721/12, 6 \cdot 12, 6595.46, 918896)$

³⁹ We also arrive at the same answer if we evaluate the incremental savings of the new loan versus the old within each period. The old-loan minus new-loan monthly payment is \$7,337.65 – \$6,595.46 = \$742.19. The old-loan minus new-loan balloon payment after six years is \$847,247 – \$918,896 = \$41,649. The present value of these incremental cash flows is \$16,220:

$$16,220 = \sum_{t=1}^{72} \frac{\$742.19}{[1 + (0.0721/12)]^t} + \frac{\$41,649}{[1 + (0.0721/12)]^{72}}$$

of the amount of capital required).⁴⁰ Then the NPV of paying off the old loan would be $\$16,220 - \$10,000 = \$6,220$. However, something important has been left out of this analysis. We have ignored the value to the borrower of the prepayment option in the old loan. The DCF-based valuation of the old loan liability, $PV(CF^{OLD})$, did not account for the positive value of this option as an asset to the borrower.

The prepayment option has a positive value to the borrower can be seen in the previous calculations. We have determined that by exercising this option today, the borrower could increase the market value of its net wealth by \$6,220, even after transaction costs. Clearly, the prepayment option is worth *at least* this much. But in paying off the old loan, the borrower extinguishes this prepayment option. An option no longer exists after it is exercised.⁴¹ The loss of this option is therefore a cost to the borrower if they prepay the old loan. How much is this option worth?⁴²

To develop some basic intuition about the nature and value of the prepayment option, let's extend our previous example. In that example, current market interest rates are 7 percent. But what will interest rates be a year from now? We don't know for sure. But suppose we can characterize the future of interest rates 1 year from now by the following probabilities:⁴³

*5% with 50% probability
9% with 50% probability*

We can calculate what the NPV of paying off the old loan would be under each of these two future interest rate scenarios. This is done simply by repeating the traditional DCF valuation procedure described in the previous section, only with the old loan advanced one more year into the future, and using the future interest rate scenarios as the basis for determining the OCC in each case. With five years paid off on the old loan, its outstanding balance would be \$950,699, and its required liquidating payment including prepayment penalty would be \$969,713. Under the 5 percent market interest rate scenario, the yield on a new loan would be 5.24 percent (with five years' maturity instead of six, and a 30-year amortization rate). Applying this discount rate to the old loan's five years of remaining cash flows to maturity gives a $PV(CF^{OLD})$ value of \$1,062,160. This results in an NPV of loan prepayment of $\$1,062,160 - \$969,713 = \$92,448$ prior to transaction costs, or \$82,448 net of these costs, under the 5 percent interest rate scenario.

Similar calculations reveal that under the 9 percent interest scenario, prepaying the old loan would have a *negative* NPV to the borrower of \$75,078 (including the \$10,000 transaction cost). With market interest rates at 9 percent, it would obviously not make sense for the borrower to pay off an 8 percent loan. However, a basic characteristic of options is that they provide their owner with a *right without obligation*. The borrower does not have to prepay the old loan just because he has the option to do so. Obviously, if the 9 percent interest rate scenario occurred, the borrower would not exercise the prepayment option on the 8 percent loan.

⁴⁰ These transaction costs would typically include third-party fees such as appraisal and title insurance costs that might not be included in the loan origination fee, or investment banker fees (in the case of equity capital). In addition, the borrower should consider his or her own costs involved in searching for replacement capital.

⁴¹ It is true that a new loan used to replace the old loan might also have a prepayment option. But this option will have very little current value because the new loan is at the market interest rate. Furthermore, the borrower would have to pay for a prepayment option in a new loan (the presence of such an option is one reason the yield on a new loan would be as high as it is), whereas the borrower already owns the option in the old loan.

⁴² As noted in Chapter 19, the prepayment option is a call option on a bond-like asset whose value is $PV(CF^{OLD})$ and whose exercise price is OLB^{OLD} .

⁴³ Historically, the annual standard deviation of the yield on long-term government bonds has been on the order of 200 basis points, the standard deviation our interest rate probability function assumed here. Thus, this numerical example, though crude, is not out of the ballpark for realistic implications.

Considering the previous analysis, we can see that the prepayment option in the old loan provides the borrower with the following **contingent values** 1 year from now:

\$82,448 in the 5% scenario
\$0 in the 9% scenario

These values are contingent because they depend on which interest rate scenario actually occurs. However, given the 50 percent probability for each scenario, we can compute the expected value, as of today, of the prepayment option next year. It is⁴⁴

$$(50\%) \times (\$82,448) + (50\%) \times (0) = \$41,224$$

How much is this expected future value worth today? To find out, we must discount it to present value using a risk-adjusted discount rate that reflects the amount of risk in the prepayment option value. Such options tend to be quite risky. Witness the large range in possible value outcomes next year between the two scenarios: \$82,448 versus nothing. Thus, we should use a high discount rate. Suppose, for example, the market would require a 30 percent expected return for investing in such an option.⁴⁵ Then the present value of the prepayment option today (given our future interest rate scenario) is \$31,711 (\$41,224/1.30).

If the prepayment option is worth anything like this much, then clearly it would not make sense to pay off the old mortgage today. Prepayment today, when interest rates are 7 percent, has an NPV of \$6,220, as we have seen, ignoring the option value. Not prepaying today, but rather waiting until next year to see what happens with interest rates, has an NPV of \$31,711, ignoring the current payoff value. These two possibilities are mutually exclusive. We cannot both pay off the old loan today and wait until next year and see if we prefer to pay it off next year instead. Each course of action rules out the other. Recall the basic NPV investment decision criterion: *Maximize the NPV across all mutually exclusive alternatives*. Thus, the classical NPV criterion tells us that we should choose the wait-and-see alternative, given the interest rate scenario we have assumed here. The \$31,711 NPV is preferable to the \$6,220 NPV.

We can also obtain this same result by properly including the prepayment option value in the valuation of the old loan. The prepayment option value is positive to the borrower, but negative to the lender.⁴⁶ Labeling $C(Prepay)$ as the market value of the prepayment option, and $D(Old)$ as the market value of the old loan, we have

$$D(Old) = PV(CF^{old}) - C(Prepay) \quad (1)$$

The market-value-based NPV of old loan prepayment, from the borrower's perspective, is therefore:

$$NPV(Prepay) = D(Old) - OLB^{old} - TC \quad (2)$$

⁴⁴ Actually, the expected value of the option next year is at least this great, as we have ignored any possibility that the option might be even more valuable if it were not exercised in the 5 percent scenario.

⁴⁵ Rigorous option value theory uses arbitrage arguments to avoid having to estimate this discount rate. However, it is not uncommon for call options to provide expected returns well in excess of 20 percent even when they are "in the money."

⁴⁶ The holder of the short position in the loan holds the long position in the prepayment option. Similarly, the holder of the long position in the loan holds the short position in the prepayment option.

where TC is the borrower's transaction costs in the deal.

Equation (2) simply says that when the market value of the existing loan exceeds the cash that would be required to pay off that loan (including the transaction costs), then refinancing will have a positive impact on the borrower's net wealth. As the market value of the old loan already incorporates the value of the prepayment option (which reflects the possible value of waiting to prepay the old loan later), paying off the old loan will be currently optimal for the borrower whenever equation (2) is positive. Combining (1) and (2), we see that

$$NPV(Prepay) = PV(CF^{old}) - C(Prepay) - OLB^{old} - TC \quad (3)$$

This makes it clear that the $C(Prepay)$ component is left out of the traditional refinancing analysis. Because option value is always positive, the traditional, purely DCF-based approach will be biased in favor of refinancing, tending to give too high a value for the NPV in equation (3). In fact, because of the value of the prepayment option, it will never be optimal to pay off a loan as soon as market yields drop just a little bit below the interest rate on the old loan.

It is important to recognize that in a highly liquid debt market, $D(Old)$ could be observed empirically from the prices (yields) of traded bonds or mortgage-backed securities, and these prices would include the value of the $C(Prepay)$ component. However, most commercial mortgages are unique and held privately as whole loans, not securitized or traded on the bond market. It is often difficult to find a liquid asset closely comparable to a given mortgage.⁴⁷ This makes it necessary to evaluate the prepayment option component explicitly in order to obtain a precise computation of the market-value-based NPV of loan prepayment.

The fully rigorous evaluation of this option is a pretty technical affair, requiring the use of option theory. The fully rigorous application of option theory to mortgages is more complex than what can be handled by the model we presented in Chapter 16. This is because a mortgage valuation model that considers both the “call” option to prepay and the “put” option to default (that was described in Chapter 19) must consider two independent underlying “state variables”: the bond-like asset that is the remaining value of the subject mortgage, and the equity-like value of the real estate collateral securing the mortgage. The fundamental model to address this problem was developed by Kau, Keenan, Muller, and Epperson in a 1992 paper in the *Journal of Money, Credit & Banking*.⁴⁸

If you don't want to get that technical, in practice a simple rule of thumb is often applied to deal with the prepayment option value effect. When yields on new loans get to around 200 basis points below the interest rate on the old loan, prepayment likely makes sense, even considering the prepayment option value. However, if not much maturity is left on the old loan, or if one has good reason to believe that market interest rates may fall further in the near future, then it may still not make sense to pay off the old loan.⁴⁹ On the other hand, if the loan has a long time remaining before maturity, or if one has good reason to believe that interest rates are about as low as they will go, then it may make sense to refinance, even if current interest rates are less than 200 basis points below the old loan rate.

⁴⁷ What is needed is comparability in the contract interest rate, maturity, prepayment ability (including penalties), and default risk. Similar size of the loan may also be important if there are economies of scale in prepayment transaction costs.

⁴⁸ Hilliard, Kau, and Slawson (1998) present a useful simplified approach based on an expanded version of the “binomial model” that we presented in Chapter 16.

⁴⁹ Also, a substantial prepayment penalty or abnormally large transaction costs can eliminate the value of prepayment to the borrower.

20.4 CHAPTER SUMMARY

Building on Chapters 2 and 19, this chapter presented the basic nuts and bolts for quantifying mortgage cash flows and yields. We saw how to apply this knowledge in loan valuation and in addressing related decisions such as refinancing. The procedures and methods presented here are as relevant for residential loans as for commercial loans, and as important to borrowers as they are to lenders.

KEY TERMS

- Four basic rules of payments and balances
- Outstanding loan balance (OLB)
- Contract principal (L)
- Interest owed (INT)
- Amortization of principal (AMORT)
- Payment amount (PMT)
- Contract interest rate (r) interest-only loan constant-amortization mortgage (CAM)
- Constant-payment mortgage (CPM)
- Balloon payment
- Graduated payment mortgage (GPM)
- Negative amortization
- Adjustable rate mortgage (ARM)
- Yield curve
- Margin (in ARM)
- Index (in ARM)
- Teaser rate
- Fixed-rate mortgage (FRM) yield (IRR)
- Yield-to-maturity (YTM) origination fee discount points basis points
- Par value
- Annual percentage rate (APR)
- Effective interest rate prepayment penalty mortgage menu
- Mortgage-equivalent yield (MEY)
- Bond- or coupon-equivalent yield (BEY, CEY)
- Prepayment option
- Refinancing
- NPV of refinancing
- Contingent value

STUDY QUESTIONS

Conceptual Questions

- 20.1 Describe the four basic rules for computing loan payments and balances.
- 20.2 What are the major advantages and disadvantages of interest-only loans for both borrowers and lenders, as compared to amortizing loans?
- 20.3 In what type of economy might a CAM be most useful? What are the major problems with CAMs in a non-deflationary economy?
- 20.4 What advantages do CPMs have over CAMs and interest-only loans?
- 20.5 Describe the general relationship among the amortization rate, maturity, regular payment, and balloon payment in a CPM.
- 20.6 What are the major advantages and disadvantages of ARMs from the borrower's and lender's perspective? How does an ARM subject the lender to less interest rate risk than

does a CPM? In what economic circumstances (e.g., interest rate environment) will ARMs be most useful?

20.7 What are the conditions in which the YTM will equal the contract interest rate on the mortgage?

20.8 At the time of loan issuance in the primary market, what can we say in general about the relationship between the YTM and the contract interest rate if the loan has disbursement discount points?

20.9

- What is the difference between the quoted YTM and the expected return (going-in IRR) for the typical mortgage investor? (Hint: Describe two sources of difference.)
- When will the expected return (over the realistic prepayment horizon) exceed the YTM in a default-free loan?

20.10 What is the APR, and how is it related to the YTM?

20.11

- What are some of the major reasons up-front points and fees are so common in the mortgage business?
- What is the major reason for the existence of prepayment penalties?

20.12

- How can the refinancing decision be evaluated as a comparison of two loan values?
- How must you define the new loan (or replacement loan) in such an analysis?
- Why is the refinancing decision not, fundamentally, a comparison of two loan values? What is it instead?

20.13 What is the relevant opportunity cost of capital to use as the discount rate in a refinancing evaluation?

20.14 Why is it important to keep the loan amount constant between the old and new loan in a refinancing analysis?

20.15 What is left out in the traditional DCF-based valuation of the refinancing decision?

20.16

- Why might it not make sense to refinance a loan with a prepayment option as soon as the market yield on new loans dips below the contract rate on the old loan?
- How is the prepayment option value incorporated in the classical NPV investment decision criterion applied to the refinancing decision?

20.17 Describe the rule of thumb commonly applied in practice to deal with the effect of the prepayment option on the refinancing decision.

Quantitative Problems

20.18 Consider a \$2 million, 8%, 30-year mortgage with monthly payments. Compute the first three payments and the loan balance after the third payment for each of the following loan types: (a) interest-only, (b) CAM, (c) CPM.

20.19 Compute the mortgage constants for the interest only and CPM mortgages in 20.19. (The annual mortgage constant (MC) is the annual debt service per dollar loan amount, or $MC = (12 * PMT) / LOAN$.)

20.20 Consider a \$2 million, 30-year amortization ARM with monthly payments and annual interest adjustments. The initial interest rate is 6%. The index for the loan is 1-year U.S. Treasury bonds, currently yielding 5.5%. The loan has a margin of 250 basis points.

- Is the loan's initial interest rate a teaser rate? How do you know?
- If 1-year T-bonds remain at 5.5%, what will be the applicable interest rate for this mortgage after the first year?
- What are the initial monthly payments on this loan?

- d. Assuming T-bonds remain at 5.5%, what will be the monthly payments after the first year?
- e. Under that assumption (and assuming no discount points), what is the forecasted yield-to-maturity on this loan at the time it is issued, assuming it has no discount points?

20.21 Consider a \$120,000, 25-year amortization ARM with monthly payments and annual interest rate adjustments. The initial interest rate is 6.5%. The index for the loan is one-year U.S. Treasury bonds, currently yielding 4.75%. The loan has a margin of 250 basis points, and *interest rate caps* of 2% per year at each adjustment and 4% over the life of the loan. It is expected that 1-year bond yields will increase to 6.25% over the second year (at the end of year 1) and to 8% during year 3 (at the end of year 2).

- a. Compute the payments and balances over the first three years. (Hint: With interest rate caps in effect, the contract rates in year 2 and beyond equal the minimum of the following three quantities: (i) index + margin (i.e., rate in the absence of caps); (ii) previous period rate + annual adjustment cap; (iii) initial period rate + life-of-loan cap.)
- b. Assuming a loan origination fee of 2 points paid up front, determine the yield (effective cost) assuming a three-year holding period. (Hint: Use your “cash flow” keys and solve for the IRR.)

20.22 Consider a \$2 million, 8% CPM with monthly payments. What are the regular monthly payments and the balloon payment amounts in each of the following cases?

- a. Fully amortizing, 25-year loan
- b. 25-year amortization, 10-year balloon
- c. 15-year amortization, 10-year balloon
- d. What are the major disadvantages and advantages of the 15-year amortization-rate 10-year loan in part (c) as compared to the 25-year amortization-rate 10-year loan in part (b)?

20.23 Consider a \$2 million, 7.5%, 30-year mortgage with monthly payments. What is the YTM of this loan under the following circumstances?

- a. No points, fully amortizing
- b. Two points of disbursement discount, fully amortizing
- c. Two points of disbursement discount, eight-year maturity with balloon

20.24 Consider a \$2 million, 7.5%, 30-year mortgage with monthly payments and an expected realistic prepayment horizon of eight years. What is the contractual yield (effective interest rate) at issuance over the expected life of the loan under the following circumstances?

- a. No points or penalties
- b. One point of disbursement discount
- c. Two points of disbursement discount
- d. Two points of disbursement discount plus one point of prepayment penalty

20.25 Consider a \$2 million, 7.5%, 30-year mortgage with monthly payments and an eight-year maturity with balloon.

- a. How much is this loan worth at issuance if the market YTM for such loans is 7.125% BEY?
- b. If instead the market yield is 7.875%, how many disbursement discount points must the lender charge to avoid doing a negative NPV deal?

20.26 A lender wants to achieve a 7.5% yield (MEY) on a 30-year amortization, monthly-payment loan with an 8-year maturity with balloon. How many disbursement discount points must the lender charge under the following circumstances?

- a. Contract interest rate is 7.25%.
- b. Contract interest rate is 7.0%.

20.27 a. As a borrower, which of the following two 25-year, monthly-payment loans would you choose if you had a 15-year expected prepayment horizon: 6% interest rate with four points, or 6.75% interest with one-half point? b. Suppose your prepayment horizon was five years.

20.28 You work for an investment company that is considering the purchase of a “seasoned” mortgage loan that was originated two years ago. The original principal amount of the loan was \$2 million, and it has an interest rate of 8%, a 15-year term and amortization period, monthly payments, and no prepayment penalty. You decide to value the loan based on the assumption that the loan is held to full maturity (i.e., there is no prepayment). Since interest rates have fallen somewhat in the past year, you offer a price based upon a required return or yield to maturity of 7.25%.

- What price do you offer for the loan package (i.e., under your assumptions, what is the current market value of the loan)? Compare this to the current OLB or “book value” of the loan, and explain the difference.
- Assume you purchase the loan at the price you determined in part (a) and that interest rates continue to fall afterwards, and as a result, the borrower prepays the mortgage 14 months from today. What yield or IRR will you have realized?
- Consider the situation in part (b). Determine the size of the prepayment penalty (\$ amount) that would have ensured that you earned a yield (or IRR) of 7.25%.

20.29 Five years ago, you took out a \$120,000, 30-year mortgage at 9% interest, with no prepayment penalty. Today you can get a 25-year mortgage at 7% with 2 points in closing costs. Assuming you plan to be in your house for another five years, determine the NPV of refinancing (ignoring option value considerations). Do the calculation twice, once using the traditional refinancing calculation and then again using the shortcut (or comparison of OLB and market value approach). Of course, your answer should be the same for each.

20.30 (Excel Exercise) An insurance company is negotiating contract terms on a commercial mortgage loan with a potential borrower. The loan in question is a monthly payment, 20-year amortization, fixed-rate balloon loan with a 7-year term (i.e., maturity), an annual interest rate of 5.85% and 1.5 points of disbursement discount.

- The yield to the lender assuming the loan is held to the end of the term (assume a loan amount of \$100). You should do this first with your financial calculator and then replicate your calculations in Excel; create a spreadsheet that takes the mortgage contract terms (amount, rate, etc.) as inputs and then use financial functions [=PMT(), =FV(), etc.] to calculate PMT and OLB. Show the mortgage cash flows from month zero to month 84 in a single column, and use the =IRR() function to solve for the yield to the lender.
- Now suppose the borrower wants to modify the loan cash flows as follows: payments are interest-only for the first two years in the life of the loan, and then amortized payments start in the first month of the third year, with a 20-year amortization period (hence, payments in years 3 through 7 are the same as those you calculated in the original loan in part (a)). The lender agrees but only if the borrower promises to make a lump sum payment, in addition to the OLB at the end of year 7, such that the lender's yield (IRR) is 6.5%. Determine the size of the lump sum payment. You should do this in the Excel spreadsheet you created for part (a). Specifically, add another column, make the required changes to the PMTs and OLB(84), and determine the additional lump sum required to make the IRR equal 6.5%.

*20.31 (Prepayment Option) Three years ago, you obtained a 10%, \$6 million, monthly-payment mortgage with 20-year amortization and an 8-year maturity. (The loan thus matures five years from now, with a balloon payment.) This loan has a prepayment clause, but

stipulates a three-point prepayment penalty on the outstanding balance. Today, it would be possible to obtain a similar mortgage at 8% interest with a one-point origination fee up front and 20-year amortization.

- a. Assuming transaction costs would be \$60,000, and the current value of the prepayment option in the old loan is \$150,000, what is your NPV for paying off the old loan today?
- b. If you could reduce the transaction costs to \$50,000, should you pay the loan off immediately?

*20.32 (Wraparound Loan) Consider the same old loan as in Question 17.36, only now suppose interest rates have risen instead of fallen, so that similar loans today would carry a 12% interest rate. Suppose further that the old loan has no due-on-sale clause, and you want to sell the property that is collateral for the loan. A buyer is willing to pay \$10 million for the property, but has only \$2 million available, and does not feel comfortable with the payments on an \$8 million mortgage at 12%, although he would do the deal at 11.5% with a 30-year amortization rate and a 5-year balloon.

- a. What would be your yield on your investment in a wraparound loan meeting the seller's specifications?
- b. Why are you able to get an expected return on this investment so much in excess of the current market rate of 12%?

*20.33 (Yield Maintenance) A commercial mortgage is written for \$1 million at 8% with 30-year amortization and a 10-year balloon payment. A yield-maintenance (YM) prepayment penalty is included as follows. If the borrower pays the loan off early, she must pay the lender an amount such that if the lender reinvests the proceeds (including the prepayment penalty) in U.S. Treasury bonds maturing on the same date as the original maturity of the mortgage, the lender will receive the same 8% mortgage-equivalent YTM on the loan's outstanding balance as she would have received in the mortgage over the remaining time until the loan's original maturity. Now suppose the borrower prepays the loan after seven years. Suppose that on that date, 3-year government bonds are yielding 6% (bond-equivalent yield). How much prepayment penalty must the borrower pay? Compute your answer based on coupon-equivalent yield, converting the mortgage MEY to its equivalent BEY.

*20.34 (Yield Maintenance) A commercial mortgage is written for \$1 million at 8% with 30-year amortization and a 10-year balloon payment. A yield-maintenance prepayment penalty is included as follows. If the borrower pays the loan off early, he must pay the lender the present value of the remaining contractual payments on the loan, discounted at the then-prevailing rates on T-bonds of a maturity equivalent to the remaining time on the loan, plus a margin of 50 basis points. (The difference between MEY and BEY is ignored, that is, the BEY T-bond rate plus margin is applied directly, to the remaining monthly cash flows in the loan as if it were a MEY rate.) Now suppose the borrower prepays the loan after seven years. Suppose that on that date, 3-year government bonds are yielding 5.50%. How much prepayment penalty must the borrower pay?

21 Commercial Mortgage Analysis and Underwriting

CHAPTER OUTLINE

- 21.1 Expected Returns versus Stated Yields: Measuring the Impact of Default Risk
 - 21.1.1 Yield Degradation and Conditional Cash Flows
 - *21.1.2 Hazard Functions and the Timing of Default
 - 21.1.3 Yield Degradation in Typical Commercial Mortgages
- 21.2 Commercial Mortgage Underwriting
 - 21.2.1 Basic Property-Level Underwriting Criteria
 - 21.2.2 Variables and Loan Terms to Negotiate
 - 21.2.3 Numerical Example of the Normative Underwriting Process
- 21.3 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- How to quantify the effect of default risk on the expected returns to commercial mortgages.
- How commercial mortgage underwriting procedures are related to default risk.
- The major traditional procedures and measures used in commercial mortgage underwriting in the United States.

The previous two chapters familiarized you with the fundamentals of mortgages, basic concepts and considerations that are relevant to both borrowers and lenders, for both residential and commercial mortgages. In this chapter, we will focus more specifically on issues that are most relevant to commercial mortgage lenders and investors. (Of course, an appreciation of the lender's perspective can also be very useful to borrowers who must negotiate with lenders.) This chapter is divided into two main parts. After an introduction discussing some quantitative aspects of default risk, we will describe the nuts and bolts of typical commercial mortgage underwriting.

21.1 EXPECTED RETURNS VERSUS STATED YIELDS: MEASURING THE IMPACT OF DEFAULT RISK

The previous chapter described in depth how to compute the yield-to-maturity (YTM) for a mortgage. We noted that the yield computed in this way is a **contract yield**, or what is often referred to as a **stated yield**, as distinct from a realistic **expected return**. The former is based on the contractual cash flow terms of the mortgage, while the latter recognizes the realistic probability of default and foreclosure. To distinguish these two measures of return, in the present chapter, we will use the abbreviation YTM to refer to the contract yield, and $E[r]$ to refer to the expected return (which may also be referred to as the *expected yield* or *ex-ante yield*).¹

¹ Note that the use of $E[r]$ in this context is consistent with our previous use of this label in Parts III–V of this book where it referred to the mean of the probability distribution of the future return.

In practice, quoted yields are always contract yields, and people work with contract yields when designing and evaluating mortgages. However, for mortgage *investors*, whether they are loan originators or buyers in the secondary market, realistic expected returns are the more fundamental measure for making investment decisions. The difference between the stated yield and the expected return quantifies the impact of the default risk on the ex-ante return that the lender cares about.

Let's look at a quick and simple numerical example to clarify some basic terminology. Consider a 3-year, annual payment, interest-only (non-amortizing) mortgage, with no fees or disbursement points up front. This mortgage originated for \$1 million at a 6 percent interest rate. This 6 percent rate will be the mortgage's contractual interest rate or "**coupon**" (rate) throughout the life of the loan, regardless of what happens in the bond market. The equation below describes these contractual terms on the loan:

$$\$1,000,000 = \frac{\$60,000}{1.06} + \frac{\$60,000}{1.06^2} + \frac{\$1,060,000}{1.06^3}$$

Suppose the bond market suddenly changes, and prevailing interest rates drop a whole point so that immediately after being issued, this mortgage now commands a price of \$1,027,232. The same mortgage with the same contractual cash flows would have a market yield or YTM of 5 percent, reflecting the bond market. This would be expressed in the present value DCF equation format as follows:

$$\$1,027,232 = \frac{\$60,000}{1.05} + \frac{\$60,000}{1.05^2} + \frac{\$1,060,000}{1.05^3}$$

This means that the mortgage at its current market price presents investors with a 5 percent expected return (in the form of a going-in IRR) *provided* it pays all three of its annual cash flows according to contract as indicated in the equation. But suppose in the third year the lender was not able to obtain the \$60,000 interest payment, but only recouped the \$1 million principal amount. The realized IRR (effective yield in this case) on that three-year cash flow stream would be 3.10 percent. Note that we solved for the 3.1 percent yield, which matched all the cash flows received with the initial investment of \$1,027,232:

$$\$1,027,232 = \frac{\$60,000}{1.031} + \frac{\$60,000}{1.031^2} + \frac{\$1,060,000}{1.031^3}$$

Finally, suppose that there is a 90 percent chance that the loan will pay off as contracted, including all three annual interest payments and return of principal, and a 10 percent chance that it will default in the third year as described above, with these two possibilities exhausting all possible outcomes of the loan. Then the *expected return* (what we will label $E[r]$) on the loan investment at the \$1,027,232 price, based on these two cases shown above and their corresponding probabilities, would be:

$$E[r] = (0.9)5.00\% + (0.1)3.10\% = 4.81\%$$

The math is based on the weighted probabilities times the resulting yields. The 19 basis-points difference between the 5 percent contractual YTM (on the 6 percent coupon loan) and the 4.81 percent expected return on the loan is due to the possibility of default on the loan. This difference is obviously important for investors (and indeed for borrowers as well). The sections that follow will introduce you to this consideration in more depth.

21.1.1 YIELD DEGRADATION AND CONDITIONAL CASH FLOWS

In the commercial mortgage business, shortfalls to the lender as a result of default and foreclosure are referred to as **credit losses**. The effect of credit losses on the realized yield as compared to the contractual yield is referred to as **yield degradation**. To understand how to quantify the difference between stated and expected yields in mortgage investments, let's begin by examining yield degradation with a simple numerical example, using our previous simple 3-year loan example, with some slight modifications.

Suppose the loan is for \$100 at an interest rate of 10 percent, such that the contractual cash flows call for two annual payments of \$10 followed by \$110 at the end of the third year. As there are no points or origination fees, the contractual YTM on this loan at origination is 10 percent. Therefore, at the end of the third year, the realized yield on the loan will also be 10 percent if there is no default.

Now suppose that the borrower makes his first two payments but defaults when it comes time for the third payment. The lender takes the property and sells it immediately but is only able to get 70 percent of what is owed, or \$77. In this case, the credit loss is \$33, and 70 percent is referred to as the **recovery rate**, or alternatively, a **loss severity** of 30 percent. The result would be a realized yield on the mortgage of -1.12 percent, computed as the IRR on the actual cash flows to the lender:

$$0 = -\$100 + \frac{\$10}{1+(-0.0112)} + \frac{\$10}{[1+(-0.0112)]^2} + \frac{\$77}{[1+(-0.0112)]^3}$$

The yield degradation is 11.12 percent, as the realized yield of -1.12 percent is this much less than the contract yield of 10 percent. The yield degradation is how much the lender loses compared to what she was supposed to get, measured as a multiperiod lifetime return on the original investment.

From an ex-ante perspective, analyzing the mortgage beforehand, we would refer to this 11.12 percent yield degradation as a **conditional yield degradation**. It is the yield degradation that will occur *if* the loan defaults in the third year, and *if* the recovery rate is 70 percent. The 70 percent recovery rate is also a conditional rate, conditioned on the default occurring in the third year. Thus, ex-ante, the 11.12 percent yield degradation is *conditional* on these events or assumptions.

Suppose the default occurs in the second year, with 70 percent recovery. In this case, the realized yield would be — 7.11 percent, as

$$0 = -\$100 + \frac{\$10}{1+(-0.0112)} + \frac{\$77}{[1+(-0.0112)]^2}$$

The yield degradation would therefore be 17.11 percent. This points to an important fact. Other things being equal (in particular, the conditional recovery rate), *the conditional yield degradation is greater, the earlier the default occurs in the loan life*. Lenders are hit worse when default occurs early in the life of a mortgage. (This has important implications for construction loans, which are inherently shorter-term in duration. If there is a default in a construction loan, it will be early in the lender's investment.)

Now let's consider the relation among the contract yield, the conditional yield degradation, and the expected return on the mortgage. After all, the expected return is most relevant to the mortgage investor at the time when the decision is made to invest in the loan (i.e., to issue the loan). In other words, it is an unconditional ex-ante measure, so we must specify the ex-ante probability of default and the expected conditional recovery rate or loss severity and derive the unconditional expected return from that. As in our introductory example, suppose that at the time the mortgage is issued, there is a 10 percent probability of default in the third year with 70 percent recovery and no chance of any other default event. Thus, there is a 90 percent chance the loan will return the contract yield, and a 10 percent chance it will return that yield less the third-year conditional yield degradation of 11.12 percent. Under these circumstances, and assuming these probabilities are based on past empirical experience, the expected return at the time the mortgage is issued is approximately 8.89 percent, computed as:

$$\begin{aligned} E[r] - 8.89\% &= (0.9)10.00\% + (0.1)(-1.12\%) \\ &= (0.9)10.00\% + (0.1)(10.00\% - 11.12\%) \\ &= 10.00\% - (0.1)(11.12\%) \end{aligned}$$

Thus, we can express the expected return as equal to the contract yield (labeled *YTM*), minus the product of the default probability (labeled *PrDEF*) times the conditional yield degradation (labeled *YDEGR*):

$$E[r] = YTM - (PrDEF)(YDEGR)$$

*21.1.2 HAZARD FUNCTIONS AND THE TIMING OF DEFAULT

In this example, we simplified the problem of computing the expected return by assuming there is only one point in the life of the mortgage when default can occur. In reality, default conceivably can occur at any point in time. This possibility is represented by what is known as a **hazard function**. The hazard function tells the conditional probability of default at each point in time given that default has not already occurred before then. As an example, let's return to the simple three-year loan again. Suppose that the hazard function representing the conditional probability of default in each year is given in the following table:

Year	Hazard
1	1%
2	2%
3	3%

This means that there is a 1 percent chance of default in the first year (i.e., at the time of the first payment), a 2 percent chance in the second year *if the loan has not already defaulted in the first year*, and a 3 percent chance in the third year if the loan is still in effect by then.

Once we know the hazard function for a mortgage, we can compute the cumulative and unconditional default and survival probabilities, as shown in the following table.

Year	Hazard	Conditional Survival	Cumulative Survival	Unconditional PrDEF	Cumulative PrDEF
1	0.01	$1 - 0.01 = 0.9900$	$0.99 \times 1.0000 = 0.9900$	$0.01 \times 1.0000 = 0.0100$	0.0100
2	0.02	$1 - 0.02 = 0.9800$	$0.98 \times 0.9900 = 0.9702$	$0.02 \times 0.9900 = 0.0198$	$0.0100 + 0.0198 = 0.0298$
3	0.03	$1 - 0.03 = 0.9700$	$0.97 \times 0.9702 = 0.9411$	$0.03 \times 0.9702 = 0.0291$	$0.0298 + 0.0291 = 0.0589$

The **conditional survival probability** in each year is 1 minus the hazard for that year. It is the probability that the loan will not default in that year, given that it has not defaulted prior to that year. The **cumulative survival probability** is the probability that the loan has survived (not defaulted) through the given year. This is the product of all the previous conditional survival probabilities (including the given year). The **unconditional default probability** is the probability, as of the time of loan origination, that default will occur in a given year. It equals the hazard in that year times the cumulative survival probability through the end of the previous year.² This is an unconditional probability because it does not depend on the conditioning assumption of the loan not having yet defaulted prior to the given year. The cumulative default probability is the probability, as of the time of loan origination, that the loan will have defaulted during or prior to the given year. It equals the sum of the unconditional default probabilities up to and including the given year, and it also (equivalently) equals one minus the cumulative survival through the given year.

² The unconditional default probability can also be computed as the difference in the cumulative survival probability at the beginning and end of the given year (e.g., $0.0291 = 0.9702 - 0.9411$).

Thus, the probability that the loan will default at any time during its life is the **cumulative default probability** through the end of the loan maturity, or (equivalently) the sum of all the unconditional default probabilities across all the years in the loan life. With hazards of 1 percent, 2 percent, and 3 percent, the ex-ante probability that Bob will default at some point in the life of his loan is 5.89 percent (computed as $0.0589 = 0.0100 + 0.0198 + 0.0291 = 1 - 0.9411$).³

For each year in the life of the loan, a conditional yield degradation can be computed, conditional on default occurring in that year, and given an assumption about the conditional recovery rate in that year. For example, we saw that with our previous example 10 percent loan, the conditional yield degradation was 11.12 percent if default occurs in year 3, and 17.11 percent if default occurs in year 2, in both cases assuming a 70 percent recovery rate. Similar calculations reveal that the conditional yield degradation would be 22.00 percent if default occurs in year 1 with an 80 percent recovery rate.⁴

Defaults in each year of a loan's life and no default at all in the life of the loan represent mutually exclusive events that together exhaust all of the possible default timing occurrences for any loan (analyzing at the annual frequency). For example, with the three-year loan, it will either default in year 1, year 2, year 3, or never. Thus, the expected return on the loan can be computed as the contractual yield minus the sum across all the years of the products of the *unconditional* default probabilities times the conditional yield degradations. This is simply a generalization of formula (1):

$$E[r] = YTM - \sum_{t=1}^T (PrDEF_t)(YDEGR_t)$$

where $PrDEF_t$ is the unconditional default probability in year t , $YDEGR_t$ is the conditional yield degradation in year t , and there are T years in the life of the loan.

For example, given the previously stated hazard function (1 percent, 2 percent, and 3 percent for the successive years) and conditional recovery rates (80 percent, 70 percent, and 70 percent for the successive years), the expected return on our example 10 percent mortgage at the time it is issued would be

$$\begin{aligned} E[r] &= 10.00\% - [(0.0100)(22.00\%) + (0.0198)(17.11\%) + (0.0291)(11.12\%)] \\ &= 10.00\% - 0.88\% \\ &= 9.12\% \end{aligned}$$

The 88-basis-point shortfall of the expected return below the contractual yield is the **ex-ante yield degradation** in Bob's mortgage. This might also be referred to as an **unconditional yield degradation**. It reflects the ex-ante credit loss expectation in the mortgage as of the time of its issuance.⁵

³ Note that this is less than the sum of all the individual hazards, which would be 6.00%.

⁴ This is computed as $(80\%) \times \$110/[1 + (-0.12)] = \100 , so -12% is the conditional return, and -12% is 22% less than the contractual yield of 10 percent. Note that in computing the conditional yield, the conditional recovery rate is applied to the OLB at the time of the default, in this case \$110. This reflects the fact that the lender would accelerate the loan in the event of default so that the entire outstanding balance would be due at that time, including interest owed. It is likely that a greater proportion of the loan balance could be recovered from a foreclosure in the first year, as there is less time for the property to have lost value subsequent to loan issuance. So it is plausible to assume a recovery rate of 80 percent in the first year if the subsequent rate is 70 percent.

⁵ An alternative method of computing the expected return may be employed. Instead of taking the average across the IRRs associated with each future scenario, we could compute the IRR of the average cash flow across the future scenarios. Instead the average IRR: $E[r] = E[IRR(CF)]$ as above, we could define and compute the expected return as the IRR of the average cash flow: $E[r] = IRR(E[CF])$. The values are not exactly the same, though they are usually similar (with the average cash flow method indicating a few basis points less yield degradation than the average IRR method). The average cash flow method is implicitly assumed in calculations of expected returns for property equity investments and in standard DCF applications as introduced in Chapter 4. However, the “average IRR” method described in the text here makes sense from an “investment manager’s perspective,” where investor preferences are defined on returns achieved.

21.1.3 YIELD DEGRADATION IN TYPICAL COMMERCIAL MORTGAGES

How large is the ex-ante yield degradation in typical commercial mortgages in the United States? To answer this question, we need to know something about the hazard function that is typical of commercial mortgages. Surprisingly enough, there was relatively little publicly available information on this subject until the 1990s.

Since then, a growing number of studies have contributed to our ability to analyze commercial mortgages rationally. Initially, the major data source used for the most high-profile studies in the institutional commercial property investment industry was the record of mortgages issued by members of the American Council of Life Insurers (ACLI). We noted in Chapter 19 that life insurance companies were the major traditional source of debt capital for long-term loans on larger, stabilized properties. The ACLI data goes back into the 1970s, thus presenting a long history that, importantly, included a couple of major down cycles in the commercial property market (mid-1970s, late 1980s). Another data source is the Federal Reserve Board and other government regulators of banks and depository institutions. However, banks and thrifts tend to specialize in short-term and construction loans, as well as in smaller properties many of which might be considered below the “institutional” investment threshold.⁶ A challenge was posed by the fact that a new and rather different major source of capital for commercial mortgages arose in the 1990s in the form of the CMBS industry. Of course, there was, at first, no historical database of the performance of CMBS loans. CMBS investors tended to infer from the most similar seeming source of data with a long history, that of the ACLI. But we now have a wealth of data on CMBS loan performance, most interestingly in the aftermath of the financial crisis and recession of 2008–2010.

Exhibit 21-1 portrays the empirical hazard function found in an influential study that was updated several times, initially by Snyderman, Esaki, and others at the investment banking firm of Morgan Stanley.⁷ The study examined defaults through 2002 in almost 18,000 individual commercial mortgages issued from 1972 through 1997 by major life insurance companies that are members of the ACLI.⁸ Note that the hazard function is characteristically humpbacked. The probability of default immediately after loan issuance is relatively low, but rises rapidly, peaking in the early to middle years of the loan life, before falling off in the later years. Constant-payment loans are most likely to get into trouble roughly during the third through seventh year of the life of the loan.⁹ In the ACLI data, the hazard probabilities peak at 1.8 percent in the sixth year and fall off rapidly after the seventh year. The implied mean time until default (if it occurs) is about seven years.

Exhibit 21-2 portrays the cumulative survival function implied by the hazard probabilities in Exhibit 21-1. The implied cumulative lifetime default probability is slightly under 16 percent. Almost one in six commercial mortgages issued by life insurance companies in the United States during the 1972–1997 period defaulted.¹⁰ However, as of the time of the study in the mid-2000s, a much smaller fraction of commercial mortgages issued after 1990 had defaulted. But that favorable record turned out to have been rather misleading in terms of the subsequent experience for CMBS loans.

⁶ You may want to look back at Exhibit 19-2 and the discussion thereof in the introduction to Chapter 19.

⁷ The original study was by Mark Snyderman in 1991. The reference for the version reported here is Howard Esaki, “Commercial Mortgage Defaults: 1972–2000,” *Real Estate Finance* 18(4), winter 2002. This version represented the state of the long-run empirical evidence regarding commercial mortgage default behavior as the CMBS industry went into the major boom period of the mid-2000s.

⁸ The ACLI publishes data on loan defaults in the *ACLI Quarterly Survey of Mortgage Loan Delinquencies and Foreclosures*. In the study, a loan was considered to be defaulted if it was reported as being more than 90 days delinquent. Most such loans are ultimately foreclosed or restructured or liquidated in some manner.

⁹ It would seem reasonable to expect that loans with large balloon payments, such as interest-only loans, would tend to run into trouble when they come due, although the Esaki study finds little historical empirical evidence of this.

¹⁰ Esaki et al. (1999) reported that the default rate on loans with at least ten years of seasoning (defaults through 1997 on loans issued through 1987) was 18.1 percent.

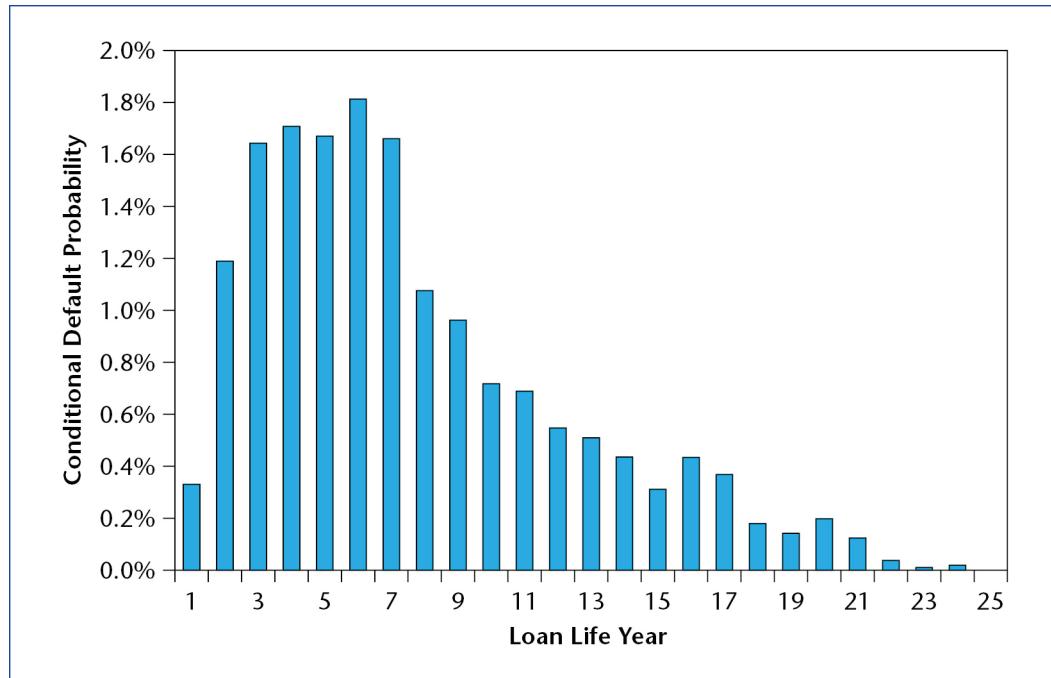


EXHIBIT 21-1 Typical Commercial Mortgage Hazard Rates.

Source: Based on data from Esaki et al. 2002.

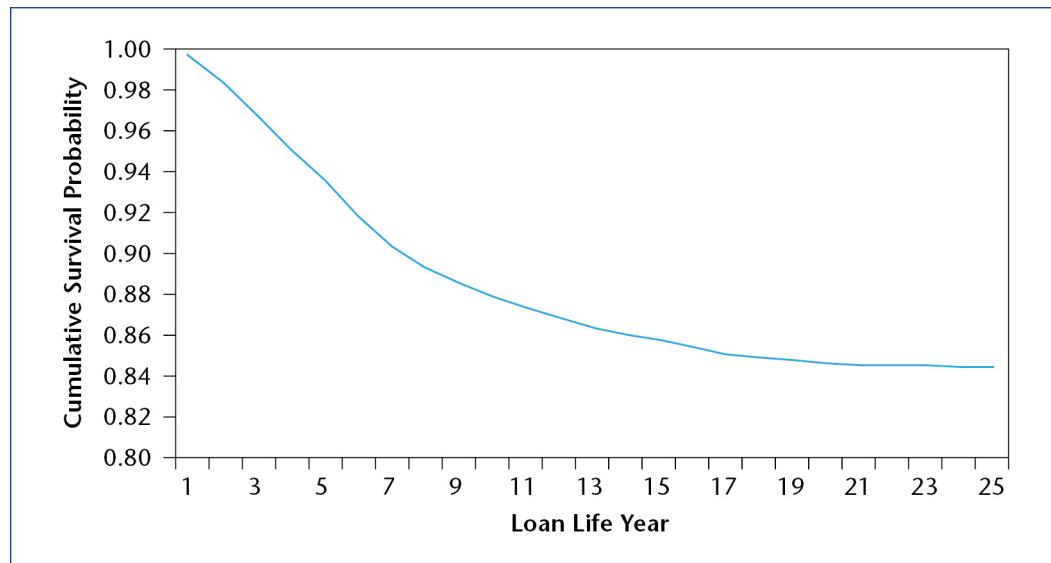


EXHIBIT 21-2 Typical Commercial Mortgage Survival Rates.

Source: Based on data from Esaki et al. 2002.

It is very important to note that historical default rates have varied widely depending on when the loans were issued. This was clear even based on the ACLI data available in the 1990s, as is depicted in Exhibit 21-3. The bars in the exhibit indicate the cumulative lifetime default rates for ACLI loans issued in each of the years 1972–1997, as compiled in the Esaki et al. study. The line

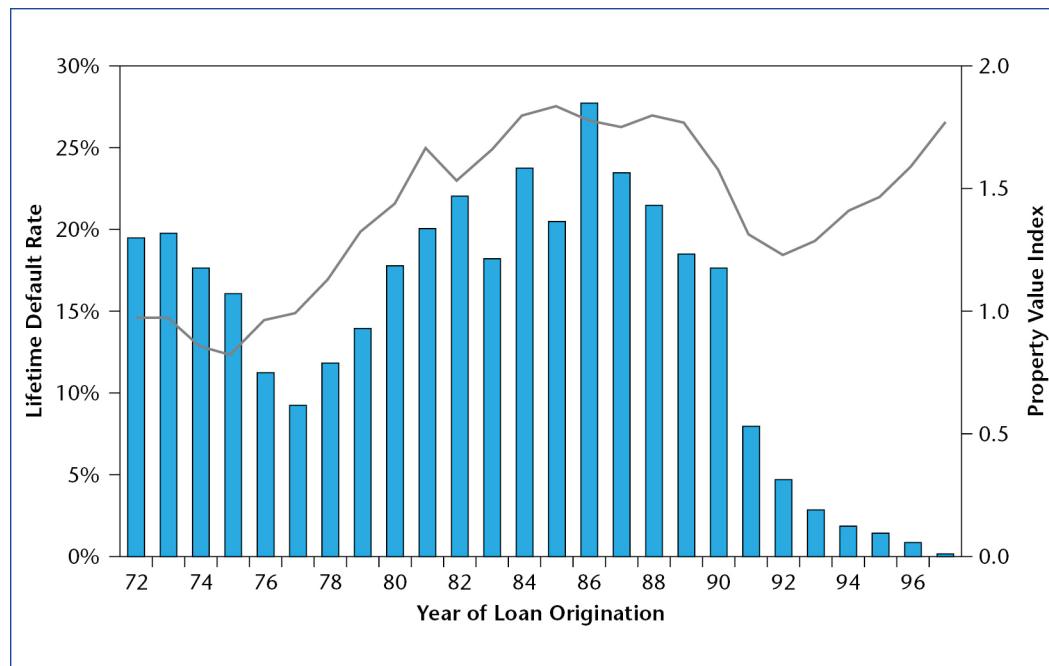


EXHIBIT 21-3 Lifetime Default Rates and Property Values.

Source: Based on data from Esaki, et al. 2002 and NCREIF Index.

in the graph traces the relative level of institutional commercial property values, based on the NCREIF index.¹¹

The historical lifetime default rates shown in Exhibit 21-3 are relatively low for cohorts of loans issued in years when property values were relatively low. For example, loans issued in the mid to late 1970s had lifetime default rates of around 10 percent. Of the loans issued in 1992, for example, less than 5 percent had defaulted by the end of 2002. On the other hand, loans issued when property values were relatively high, such as the early 1970s and mid-1980s, had much higher lifetime default rates, peaking at almost 28 percent for loans issued in 1986. When property values were relatively low, mortgage issuers granted only relatively small loan amounts, and subsequent rises in property values made default relatively rare. When property values were relatively high, larger loans were issued, and there was an increased probability that subsequent falls in property market values would put the loans “under water,” that is, the outstanding loan balance (OLB) would be greater than the value of the property collateral securing the loan.

As noted previously, this type of behavior has been described as a “Santa Claus approach,” in which borrowers are rewarded (that is, loaned more capital) when they have done well and penalized (given less capital) when they have performed poorly, during the recent past. Indeed, if real estate markets are cyclical, and the LTV ratios at which lenders will grant loans are constant, then this could result in the type of default rate cyclical suggested in Exhibit 21-3. The cycle will be only further exacerbated if lenders are actually more aggressive, as by effectively relaxing their underwriting standards, during periods when the property market is booming (and perhaps pull back to excessive conservatism just after having been “burned” in a down cycle). Of course, boom periods practically by definition correspond to times when large quantities of financial capital are flowing

¹¹ The NCREIF index appreciation value levels in Exhibit 21-3 have been “unsmoothed” by the authors and extended back in time using additional data.

into the property market, normally including a large share of debt capital in the form of commercial mortgages. Thus, the system may be structured to support boom-and-bust pattern.

The good performance on loans issued in the early 1990s was due, in part, to the fact that underwriting standards became much stricter and lending was very conservative in the immediate aftermath of the great commercial property “crash” of the late 1980s and early 1990s. Through the rest of the 1990s, lending was generally conservative. The other part of the very favorable outcome for the early 1990s loans was the strong rental market during the late 1990s, followed by the boom in property asset market values in the early 2000s. The late 1990s and early 2000s were a “forgiving” period in commercial mortgage lending.

Ironically, the very good performance of the mid-1990s cohorts of loans may have helped to lull the capital market into complacency. After all, there had been what seemed at the time a major global financial crisis in 1998, and a major stock market crash and recession in 2001–2002, and commercial mortgage performance came through those crises in pretty good shape. Many players in the industry may have been able to convince themselves that the new system, which balanced traditional “relationship” and portfolio lending from life insurance companies and commercial banks with securitized conduit lending via the CMBS industry, provided a highly efficient machine. By the mid-2000s loan underwriting became very aggressive again, especially in the new CMBS conduit segment of the market.

Exhibit 21-4 shows the history from 1990 through 2010 of the current default rate on loans outstanding by various types of lenders. The current default rate is not exactly the same thing as either the lifetime cumulative default rates by loan issuance cohorts shown in Exhibit 21-3 nor the longitudinal profile of loan lifetime hazard rates shown in Exhibit 21-1. Current default rates give a picture of how all the loans outstanding as of a given time are performing (or you can think of it as how the average such loan is performing). For example, Exhibit 21-4 indicates that in the early 1990s mortgages issued by life insurance companies were experiencing default rates of around 7 percent, as measured by the percentage of loans that were delinquent in payment by 60 days or more. If such a default rate were to continue for 10 years on all remaining loans (like an annual hazard rate), then at the end of 10 years only $(1 - 0.07)^{10} = 48\%$ of the loans would still be “alive,” a cumulative default incidence of 52 percent. But of course, such a high default rate did not persist, and we saw in Exhibit 21-3 that in fact the cumulative lifetime default incidence on the worst-performing cohort of ACLI loans (those issued in 1986) was actually less than 28 percent. But you can see how the current default rate gives a type of glimpse at the current instantaneous hazard rate facing the currently outstanding mortgages.

With this in mind, look at the picture presented by Exhibit 21-4. As bad as the previous commercial mortgage default performance was in the early 1990s in the aftermath of the commercial property down cycle at that time, the new CMBS loans were experiencing equally bad default performance by 2010 in the aftermath of the then-latest down cycle. By 2011, the delinquency rate for the worst-performing cohort of loans issued during the 2007 peak of the bubble was up to 12 percent.¹² Banks were also suffering severely by 2010, although still at¹² rates below their peak default performance in the previous cycle. Interestingly, in the down cycle of the 2000s the life insurance companies (LICs) were experiencing much less loan distress than the other major lenders (other than the two government-sponsored enterprises that were restricted to apartment loans on moderate-income-oriented multifamily rental housing). In part, it reflects differences in policies and procedures and in reporting of default rates, but this also may have been because LICs

¹² As compiled by Moody’s Investors Service “Delinquency Tracker Report” of November 2011, based on loans 60+ days delinquent or in foreclosure or REO status. By this definition the overall average CMBS loan delinquency rate in 2011 was over 9 percent.

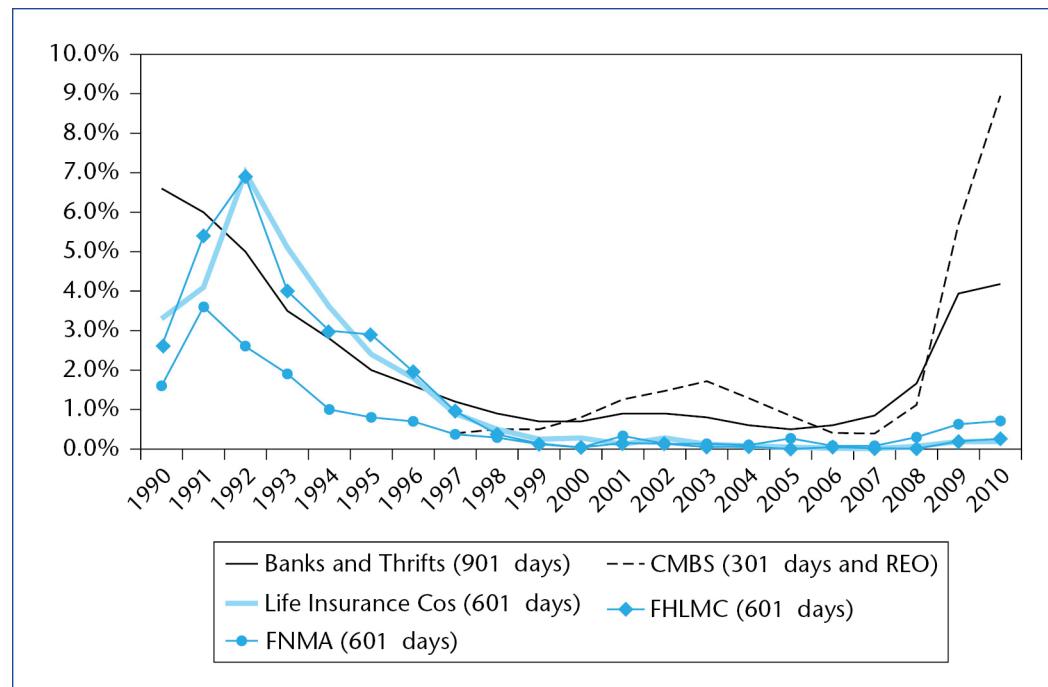


EXHIBIT 21-4 DEFAULT Rate in Outstanding Loans by Lender: 1990–2010*

*Excluding construction loans.

Source: Mortgage Bankers Association.

had learned a hard lesson in the previous down cycle and did not get quite as aggressive in their underwriting as the CMBS conduit lenders in the peak of the 2005–2007 boom. As a result, the LICs were somewhat “squeezed out” of the market during the peak of the bubble years, to their subsequent benefit.¹³

Default probabilities are not the only type of information necessary to estimate the typical magnitude of ex ante yield degradation in commercial mortgages. We also need information on conditional

¹³ The default rates shown in Exhibit 21-4 are not exactly comparable across the different types of lenders, in part because they have been defined differently in terms of the length of delinquency. For example, in the Mortgage Bankers Association report that Exhibit 21-4 is based on, CMBS default is defined by 30+ days of delinquency or REO status (real estate ownership taken over by the lender), whereas the bank and thrift default rate is defined by 90+ days of delinquency. However, there tends to be a very high correlation among delinquency duration (i.e., once a loan goes delinquent it tends often to continue to do so). It should also be recognized that the different types of lenders in Exhibit 21-4 tend to specialize in different types of commercial properties, which may therefore subject them to different severities of market conditions. For example, banks tend to focus on smaller properties and shorter-term loans (even apart from construction loans). Another reason for the “apples-versus-oranges” comparison across the different lenders in Exhibit 21-4 has to do with different policies and procedures regarding the treatment and recording of loans where the borrower is in trouble. In particular, portfolio lenders and relationship lenders (including most notably life insurance companies and banks) adopted more flexible policies with their borrowers that allowed the lenders to avoid recording loans as delinquent or nonperforming after the 2008–2010 down cycle than was the case for CMBS loans or indeed for the portfolio lenders themselves during the previous down cycle. (The new behavior became known as “pretend and extend,” and it helped the relationship lenders to avoid experiencing losses that would have been even worse.) It should also be noted that, while many loans from banks and thrifts are construction loans, which are a rather different and more risky and cyclical type of product, these have been excluded from the figures portrayed in Exhibit 21-4.

recovery rates or loss severities once default occurs. This type of information is more difficult to obtain, although more is becoming available all the time. Losses generally result to the mortgage holder even if an actual foreclosure does not happen as a result of the default. As described in Chapter 19, actions following loan default range from loan workouts to “deed in lieu” or “short sales” to formal foreclosure.¹⁴

The Esaki et al. study reported that about 55 percent of defaulting loans (defined as 90+ days delinquent in payment) were foreclosed or liquidated, 40 percent were restructured, 3 percent became delinquent again, and only 2 percent fully recovered. Among liquidated loans, the average conditional recovery rate was 69 percent (or equivalently, 31 percent loss severity). This included the effects of foreclosure expenses as well as lost interest and principal based on the *reported* value of the property at the time of foreclosure. However, a more detailed study by Ciochetti (1998) tracked 308 foreclosed mortgages from “cradle to grave,” including the period subsequent to foreclosure when the property was owned by the lender (as “real estate owned”—REO). The loans had all originated between 1974 and 1990. Ciochetti found that while the average recovery was 57 percent based on reported property value as of the time of foreclosure (which resulted in 6.5 percent conditional yield degradation), the average recovery was only 34 percent through to the final disposition of the collateral property by the lender (equivalent to 10.6 percent conditional yield degradation in his loan sample).

More recent data from Fitch Ratings provides additional insights into loss severity trends in 2023. The average loss severity (LS) on resolved loans during the year remained in line with the historical long-term average. For the \$2.1 billion in loans disposed of with losses, the overall average LS was 48.4 percent, slightly exceeding the cumulative historical average of 47.1 percent. Retail properties accounted for the majority of the losses, representing approximately 70 percent of loans resolved by balance, with an average LS of 48.2 percent. The office sector experienced the highest LS among property types, averaging 53.2 percent.¹⁵

Default probabilities and conditional recovery statistics can be combined to provide estimates of conditional and *ex ante* yield degradation for typical loans.¹⁶ Such information can be used in the construction of indices of the periodic returns to commercial mortgages, based on synthetic portfolios of typical mortgages. Such an index of commercial mortgage whole loan investment performance is the Giliberto-Levy Commercial Mortgage Performance Index (GLCMPI), which is based largely on ACLI loan performance data.¹⁷ The GLCMPI includes estimates of the magnitude of credit losses suffered in the index total return each quarter as a fraction of the OLB of all loans. On an annualized basis, these credit losses provide an approximate indication of the magnitude of the **ex-post yield degradation** realized by the aggregate portfolio of outstanding loans during each year.

¹⁴ In this context, the term “short sale” refers to a situation in which the borrower in default and the lender mutually agree to sell the property, generally due to market conditions at a price below what is owed on the loan, after which the lender takes all of the proceeds from the sale, but the lender then releases the borrower without formal foreclosure or seeking a deficiency judgment. This is similar to the “deed in lieu” except that in the latter case the borrower does not sell the property, and the lender might hold onto the property for a while as REO (real estate owned) hoping for the market to improve. This is much more possible in commercial (income producing) real estate than with owner-occupied houses.

¹⁵ Fitch Ratings, “U.S. CMBS 2023 Loan Loss Study: Retail Sector Comprises Majority of Losses,” Special Report, June 4, 2024.

¹⁶ As noted, not all defaulting loans are foreclosed, as some defaults are worked out, as described in Chapter 19. However, lenders usually still suffer losses in such cases. In addition, lenders suffer some losses from loan delinquencies that are not classified as defaults (e.g., delinquencies of less than 90 days).

¹⁷ The GLCMPI was originally published by John B. Levy & Company, Richmond, Virginia and was taken over by the Investment Property Databank (IPD) in 2011.

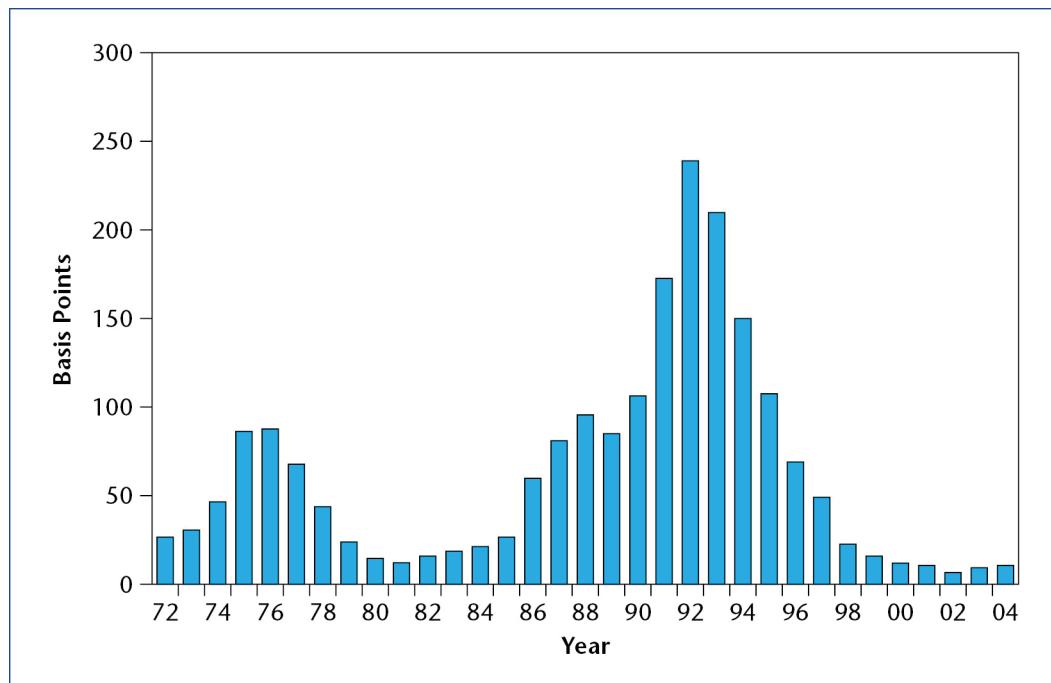


EXHIBIT 21-5 Commercial Mortgage Credit Loss as Fraction of Par Value.

Source: Based on data from GLCMPI—John B. Levy & Co.

Exhibit 21-5 shows the magnitude of annual credit losses in the GLCMPI over the 1972–2004 period (measured in basis points of outstanding par value). As we would expect, credit losses are higher during downturns in the commercial property asset market (mid-1970s, early 1990s). Credit losses peaked at almost 240 basis points in 1992, but were as low as 6–9 basis points during the boom years of the early 2000s, only to rise to close to 100 basis points again by 2010. Ex-post realizations of credit loss are certainly more volatile than expectations beforehand. However, averaged across a long span of time, ex-post credit losses provide some indication of the typical ex ante yield degradation in permanent commercial mortgages in the United States. The average annual credit loss in the GLCMPI during the 1972–2004 period was 62 basis points.¹⁸ In Exhibit 21-6, we show the default rates on commercial mortgages through 2024, and over this history from 1991 through the end of 2024 the average default rate is 3.1 percent. Losses would not be that different from the earlier history. Construction loans would be expected to contain greater ex ante yield degradation than permanent loans because of their short maturities and likely greater conditional loss severities.

¹⁸ The GLCMPI also estimates periodic total returns on commercial mortgages both with and without adjustment for credit loss. The average per annum adjusted return over the 1972–2004 period was about 90 basis points lower than the average unadjusted return. This number is greater than the 62 basis-point average credit loss because the denominator in the holding-period returns (HPRs) adjusts over time to reflect changes in the current market value of the mortgage. This value declines prior to loan default, as the current loan-to-value ratio typically rises and debt service coverage ratio typically falls, increasing the loan's default risk and thereby lowering its market value. In contrast, the denominator in the credit loss calculation remains at the par value of the loan (its contractual OLB). For this reason, in principle, the average credit loss over time probably provides a better indication of the magnitude of ex-ante yield degradation.

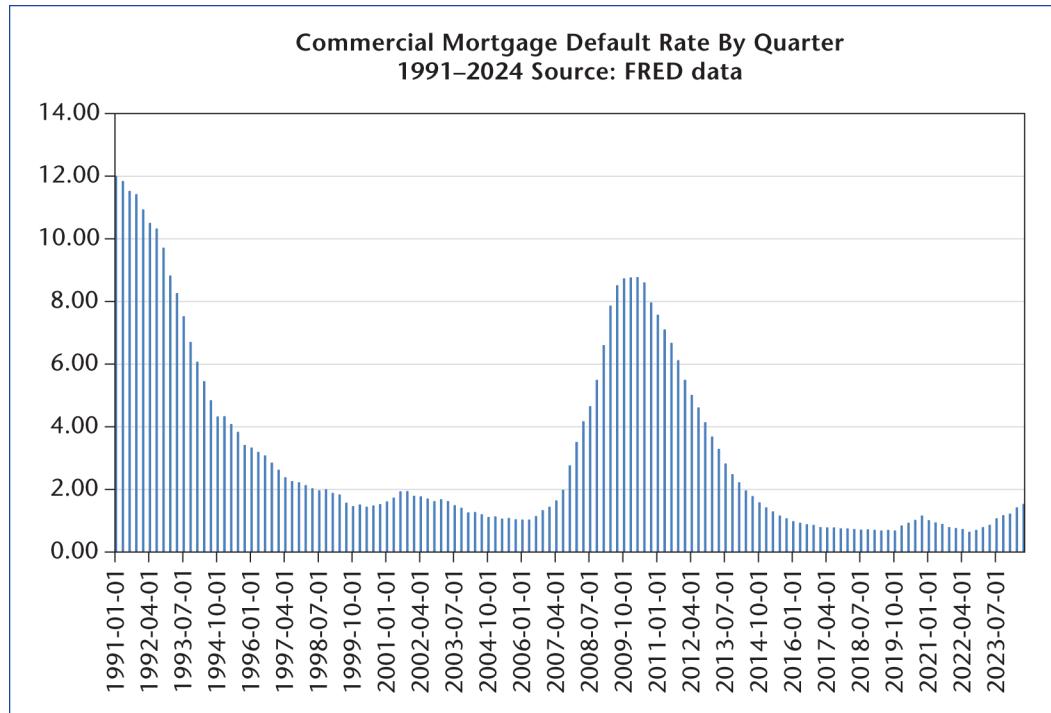


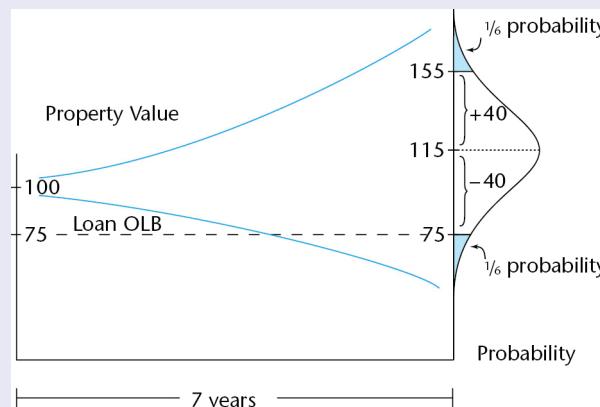
EXHIBIT 21-6 Commercial Mortgage Default Rates Through 2024.

IS IT SURPRISING THAT SO MANY COMMERCIAL MORTGAGES DEFAULT?

No borrower plans to default on a loan when he or she takes it out. Lenders seem to go to lots of trouble to try to avoid making loans that will default. So, is it surprising that about one out of every six long-term commercial mortgages in the United States defaults? Perhaps not, when you go back to fundamentals.

One of those fundamentals is *volatility* in the property market. Recall in Chapter 3 we suggested that the typical individual commercial property probably has an annual volatility in excess of 15 percent (and this doesn't include "deal noise" as described in the supplemental materials of Chapter 14). This means that the annual standard deviation in a typical property's appreciation return (or change in value) is at least about 15 percent of its initial value, including market risk and idiosyncratic risk. Now recall that the normal probability distribution has about one-third of its total probability beyond one standard deviation from its mean. This implies that about one-sixth of the probable value outcomes are at or below one standard deviation *below* the mean outcome. Suppose the expected (i.e., mean) property value a year from now is the same as the current property value. Then 15 percent annual volatility implies that there is about a one-sixth chance that a given property will be worth, one year from now, 85 percent or less of its current value. Now suppose property values follow a "random walk" through time. In other words, suppose property values change randomly across time, like stock market values. (Whether the value rose or fell during the past year does not tell you anything about whether it will rise or fall next year, like the flipping of a coin.) Then it is a mathematical fact that volatility grows with the square root of the time over which it is measured. The biennial volatility would be $\sqrt{2}$ times the annual volatility; the decennial volatility would be $\sqrt{10}$ times the annual volatility, and so forth. *

Now we know from the Esaki et al. study that the typical commercial mortgage default occurs about seven years into the life of the loan. Thus (somewhat simplistically speaking), the relevant volatility for determining default frequency is the 7-year volatility, not the annual volatility. If the annual volatility is 15 percent, then the 7-year volatility is $\sqrt{7}$ times this amount, or 40 percent. Without any expected appreciation in the property value, there would be about a one-sixth chance that a given property would be worth, seven years from now, only 60 percent of its current value, or less. Even if the property is expected to appreciate at, say, 2 percent per year, this will increase the expected value in seven years only to about 15 percent above the current value, so there would still be a one-sixth chance that the property would be worth only 75 percent of its current value, or less, after seven years. The standard loan- to-value (LTV) ratio required by mortgage issuers has traditionally been 75 percent. If loans don't amortize much and borrowers tend to default whenever the loan is "underwater" (that is, when the property is worth less than the loan balance), then the 75 percent LTV criterion suggests a typical default probability of about one-sixth. (See illustration above.)



Of course, this analysis is highly simplified, and many commercial mortgages amortize their principal over time. But, on the other hand, in a low-inflation environment, many properties may tend to appreciate at less than 2 percent per year, and we are not considering "deal noise" here (possible mispricing as of any given point in time). And while it may not make sense to default on a loan just because its par value is underwater, the reason the property has lost value is likely to be because its rental income has dropped, which may force the borrower to default. You can see the basic point: *Empirically observed default rates in commercial mortgages are consistent with empirical evidence about property volatility.*^{**} The average property may not be as risky as the average stock (whose annual volatility is typically around 30 percent), but when you lend 75 percent of a property's value, you are asking for, well, about a one-sixth chance of running into a default at some point in the life of the loan).

*If property values do not follow a purely random walk, but rather have some inertia, as would be the case in a sluggish or informationally inefficient asset market, then the volatility actually increases *more than* the square root of the measurement time interval. The analysis here would be conservative in that regard, suggesting an even greater chance of default.

^{**}Ciochetti and Vandell (1999) studied this question much more rigorously. Their estimate of the individual property annual volatility implied by the empirical evidence on commercial mortgage values and default rates is about 17 percent, assuming property values follow a random walk.

21.2 COMMERCIAL MORTGAGE UNDERWRITING

Underwriting is the word used to describe the process commercial mortgage originators go through to decide whether to issue a proposed mortgage. This process can be much more involved than taking out a typical home mortgage. Commercial mortgages are often negotiated, and the terms of the loan may be customized to reflect the unique circumstances presented by the borrower, the lender, and the collateral property. This is less true with CMBS conduit loans, which are necessarily more standardized and commoditized. But even with conduits there is often some room to maneuver or to choose among competing originators. The larger the loan is, the more it is worthwhile to spend time and energy on this negotiation and customization process. For typical commercial mortgages of several million dollars, procedures are generally fairly standardized. Even for larger loans, in the tens of millions of dollars each, standard underwriting criteria are issued by the institutions supplying the mortgage capital.

In this section, we will describe typical commercial underwriting criteria and how the underwriting process works in a normative sense. That is, our focus and perspective is on how the process is *supposed* to work. But keep in mind that, as we discussed in the preceding section, there may be times when market pressures cause the system, in effect, to “bend” the rules, even if on paper the rules may not appear to change that much. Even then, the process still utilizes the same metrics that we will discuss in this section.

What Underwriting Is All About. To begin, you should recognize the normative purpose of underwriting. In principle, the purpose of underwriting is *to make default a rare event*. We noted in the previous section that, on average, about one in six long-term commercial mortgages in the United States defaults, or at least did so during the last quarter of the twentieth century. Is this one-sixth rate “rare enough”? Well, it could be, provided lenders are being adequately compensated for the risk. As with all investments, such compensation occurs in the magnitude of the ex-ante total returns, or yields, that the lenders can realistically expect when they issue the loans. Thus, more broadly and fundamentally, underwriting has as its purpose to ensure that lenders are getting the expected return, they want at the time they make the loan. The tighter the underwriting criteria are, the lower will be the probability of default and the lower will be the ex-ante yield degradation, thereby raising expected returns toward the contractual yield rates.

Of course, lenders cannot operate outside the market. Suppose a lender tried to set underwriting guidelines so tight that the ex-ante yield degradation would be eliminated altogether. This would make the mortgage like a default-free bond. If the lender then tried to charge an interest rate above the yield on default-free government bonds, then she would probably find it difficult to make the loan. Borrowers conforming to such tight underwriting criteria could shop around among competing lenders to get a lower interest rate.

In fact, most real estate borrowers cannot (or do not want to) conform to underwriting standards so tight as to eliminate default risk. As a result, mortgage lenders have to monitor the market constantly and modify their underwriting criteria accordingly, if they want to play the game (that is, issue commercial mortgages). Loan originations must satisfy both sides of the deals, and as with any market, prices (that is, ex-ante yields) reflect an equilibrium in which, on the margin, both sides are just willing to do business. Thus, in principle, underwriting criteria ensure that the realistic expected return to the lender is high enough, but not too high, so that it includes just the ex-ante risk premium required by the market.

The Two Foci of Underwriting—Borrowers and Properties. When examining a mortgage loan application, a lender’s attention is naturally focused on two subjects: the borrower and the proposed collateral property. In the commercial mortgage business (in contrast to the residential business), the more important of these two foci is normally the property. As noted, many commercial mortgages are non-recourse, so that no personal assets are pledged beyond the property. Even when recourse is technically available beyond the collateral property to the borrower, there is often little value the lender can get. Furthermore, with a commercial mortgage (unlike a residential loan), the borrower

will rely principally on the collateral property itself, being an income property, to provide the cash to service the loan. If the property is sufficiently lucrative in comparison with the loan requirements, then the loan will probably turn out all right even if the borrower is a bit weak.

Nevertheless, the borrower is still an important consideration even in commercial mortgages. On the downside, a borrower who gets into trouble with other businesses or properties other than the collateral property may use the collateral property as a “cash cow” to bail out his other losses, perhaps to the detriment of a loan on the healthy collateral property. Borrowers of commercial mortgages can also wreak havoc on secured lenders by filing for protection under Chapter 11 of the bankruptcy law. Thus, lenders need to examine the nature of the borrowing entity and any parent or related firms and holding companies in order to ascertain the financial health of the relevant borrowing entity.

On the upside, borrowers on commercial mortgages (more so than residential borrowers) are often potential repeat customers. Commercial borrowers are often wealthy individuals, businesses, or institutions that are permanently in the real estate investment business in one way or another. If they are successful, they will need capital regularly for other projects and investments. It makes good business sense for lenders to cultivate such customers. Thus, the reputation and future business prospects of the borrower are important considerations for the lender. The mortgage lending business is very competitive in the United States, and a lender may relax underwriting criteria to some extent to cultivate or retain a borrower with a good history and solid future potential.

21.2.1 BASIC PROPERTY-LEVEL UNDERWRITING CRITERIA

Although information about the borrower is important, the primary attention when putting together the nuts and bolts of a specific deal typically focuses on several traditional criteria relating to the property that is to be the collateral in the loan. These criteria focus on two major aspects of the property: asset value and income flow. The value of the property relates naturally to the value of the loan, while the magnitude of the property’s income flow is more directly relevant to the amount of periodic debt service that will be required by the loan.

Initial Loan-to-Value Ratio (ILTV) The **initial loan-to-value ratio (ILTV)** is the classical asset-value-based underwriting criterion. It is defined as the initial loan value (the contractual principal amount of the loan) divided by the current market value of the collateral property:

$$LTV = L/V$$

The ILTV ratio is an important underwriting criterion for obvious reasons. Commercial property market values display **volatility**, which means they can go down as well as up over time.

A lower initial LTV ratio, say 50 percent LTV, will reduce the probability that at some point during the life of the loan the property will be worth less than the OLB. Although default will not necessarily occur as soon as the property is worth less than the loan balance, if the property value falls far enough below the OLB, default will certainly be rational (i.e., wealth maximizing) for the borrower. Thus, the initial LTV ratio is directly related to the ex-ante default probability of the loan. Indeed, because the current market value of the property reflects the entire future income stream the property can generate, the ILTV is arguably the most fundamental and important single underwriting criteria, as it reflects both asset value considerations (directly) and income coverage considerations (indirectly) in a single summary measure.¹⁹

¹⁹ On the other hand, the difficulty of estimating precisely the market value of the property may render the LTV a “fuzzier” measure than income-based measures, as current income can often be observed more reliably than property value. Recall in Chapter 13 we noted evidence of institutional property price dispersion around predicted values on the order of a 15% standard deviation (though valuation error for specific individual properties may be less than this).

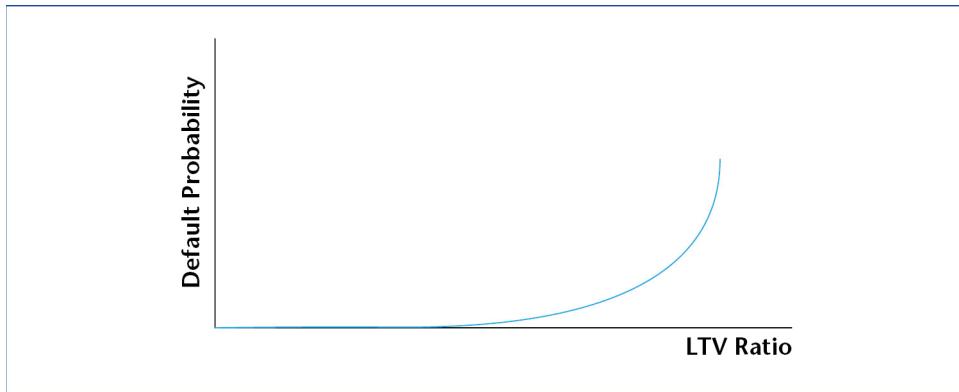


EXHIBIT 21-7 Typical Relationship between Initial LTV Ratio and the Ex Ante Lifetime Default Probability on a Commercial Property Mortgage.

For typical levels of commercial property volatility, the relationship between default probability and the initial LTV ratio on a long-term commercial mortgage looks something like the curve shown in Exhibit 21-7. Note that this relationship is highly nonlinear.²⁰ Default probability is very low and nearly constant over a wide range of low values of the LTV ratio, but then increases sharply over higher LTV ratios.

The greater the volatility in the collateral property value is, the lower will be the initial LTV ratio corresponding to a given ex-ante default probability. Typical commercial mortgage underwriting criteria in the United States have traditionally required the ILTV ratio to be equal to or below 75 percent, and sometimes no more than 65 percent. ILTV limits on property types that are viewed as riskier, such as raw land, will be set lower than this.²¹ As noted in Chapter 15, an LTV ratio of 75 percent implies a much higher leverage ratio or debt/equity ratio than is typical in the stock market. Nevertheless, during periods of rapid price inflation and during real estate booms there is often strong pressure on lenders to relax this traditional limit. The traditional limit may also be exceeded if extra guarantees and credit enhancements are included in the loan, such as extra collateral, recourse provisions, cramdown insurance, and so forth.²²

In estimating the market value of the property (V), lenders will generally require their own independent appraisal, and if the loan is to finance a property purchase, they may take the property value to be whichever is lower between the appraisal and the agreed-on purchase price of the property. In practice, the loan terms may be effectively agreed on before an independent appraisal is conducted, and lenders may often rely for decision-making purposes on their own informal estimate of property value. In any case, at least two methods of estimating property value are often employed, direct capitalization and multiyear DCF valuation (as described in Chapter 4). In keeping with principles of conservative underwriting, lenders would use the lower of the two values indicated by these two approaches for underwriting purposes. In some cases, lenders will set general guidelines governing

²⁰ In fact, this curve would typically have a hyperbolic shape very much like how we described the relationship between the “leverage ratio” and the LTV back in Chapter 12: $LR = 1/1-LTV$.

²¹ Although stated limits rarely exceed 75 percent or 80 percent, competitive pressures in the property and lending markets have likely on occasion pushed the effective (honestly stated) limit to well beyond 75 percent, with consequences in terms of subsequent default rates that were discussed in Section 21.1. Credit rating agencies rating CMBS bonds have devised more conservative and rigorous LTV metrics that during the peak bubble year of 2007 sometimes approached or exceeded 100 percent, though it should be noted that these metrics are deliberately conservative.

²² It should be noted that if property value tends to increase over time in nominal terms due to inflation, and the borrower does not add to the debt on the property, then the average LTV ratio over time will tend to be less than the initial ratio.

their in-house valuation procedures, such as placing lower bounds on the capitalization rate and/or the discount rate that can be employed in the valuation.²³

While the LTV ratio criterion is applied most stringently at the time of loan issuance, it is also often forecasted over the life of the loan based on a projection of the property NOI and direct capitalization. Of particular concern, of course, is the projected LTV ratio at the time of loan maturity if the loan is not fully amortizing. The lender needs to be able to expect that the property value will well exceed the OLB on the loan at that time.

Debt Service Coverage Ratio (DCR). As the LTV ratio is the basic asset-value-based underwriting criterion, the DCR is the classical income-based criterion. The **debt service coverage ratio (DCR)** is defined as the collateral property's annual net operating income (NOI) divided by the annual debt service required by the loan:

$$DCR = NOI/DS$$

The debt service (DS), of course, includes both the periodic interest payments and any amortization of principal called for in the loan terms. This ratio clearly makes sense as an underwriting criterion, for the NOI generated by the property is normally the primary source of cash to service the loan.

A typical standard for the DCR would require that this ratio must equal or exceed 120 percent. To reduce the risk of a cash flow squeeze, the lender wants some buffer, provided by the excess of the required DCR above 100 percent. The DCR criterion may be raised higher for types of property that appear riskier, or during times when lenders are more risk averse. On the other hand, lower DCR hurdles may be accepted during times of rapid inflation, or when the loan market is "hot." A projected temporary violation of the DCR hurdle for one or two years may be tolerated if there is a solid projection of sufficient debt service coverage during the other years of the loan life. Taken together, the LTV ratio and the DCR are the two most widely used underwriting criteria. However, other measures are also commonly examined.

Break-Even Ratio (BER) The **break-even ratio (BER)** is another widely employed income-based underwriting criterion that is usually not used instead of the DCR but as a supplemental requirement that must also be met. The BER is most simply defined as the sum of the annual debt service and property operating expenses divided by the potential gross income:

$$BER = (DS + OE)/PGI$$

The BER gives the occupancy ratio of the building (one minus the vacancy rate) below which there will be insufficient net operating income to cover the debt service. If the BER is greater than 100 percent, then the property investment will be a net cash flow drain on the borrower (even without considering any necessary capital improvement expenditures on the property). This would obviously be a dangerous situation from the lender's perspective, so underwriting criteria will typically require the BER to be less than some fraction well below 100 percent. A typical BER limit might be on the order of 85 percent or less. Sometimes the maximum BER criterion is stated as the average occupancy rate prevailing (or expected to prevail) in the space market in which the property is situated, or that rate less some buffer.²⁴

²³ For example, a lender may state that cap rates no lower than 9 percent must be employed in conducting a direct capitalization valuation. Of course, all such rigid criteria must bend to the market, as noted previously, assuming the lender really wants to be in the mortgage business. In other words, a stated standard of a 9 percent cap rate floor may look nice on paper, but if the cap rates currently prevailing in the property market are 8 percent, then the lender will not be able to issue many mortgages unless it either informally relaxes its cap rate floor or otherwise relaxes its LTV limit.

²⁴ A more complete and sophisticated formula for the BER would recognize that some expenses are essentially fixed no matter what the occupancy in the property is, while other expenses are variable. Defining "FC" as the fixed costs and

Equity-Before-Tax Cash Flow (EBTCF). As noted in Chapter 13, the **equity-before-tax cash flow (EBTCF)** is the cash flow bottom line for the equity investor on a before-tax basis, reflecting the need for capital improvement expenditures on the property. The EBTCF equals the NOI less debt service and capital improvement expenditures:

$$EBTCF = NOI - DS - CI$$

If this measure is projected to be negative for any year during the life of the loan, this raises an obvious underwriting red flag, as the borrower faces a potential negative cash flow in any such year. While the EBTCF measure is arguably more relevant in principle than the DCR, it is often relatively difficult for the lender to estimate and project future CI needs. Furthermore, CI expenditures tend to be somewhat discretionary in their amount and timing, and the property owner may be able to finance some types of CI expenditures through the use of additional debt. As a result, the DCR criterion is more widely employed in practice than the EBTCF measure, although both are often examined. The EBTCF measure is particularly relevant for certain types of property, such as properties in need of improvement and properties that employ long-term leases in which major tenant improvement expenditures (a component of CI) may be required of the landlord whenever a lease expires.²⁵

Loan Yield After the 2008–2009 financial crisis many lenders and CMBS credit rating agencies began to consider another metric, known as the “**loan yield**” (or sometimes the “**loan cap rate**”), defined as the property’s current NOI divided by the loan amount (or the remaining balance due on the loan, to update the measure as the loan amortizes). The rationale behind this metric is that it relates the loan amount to a property-pricing metric that would become relevant should the borrower default and the lender take over the property as “real estate owned” (REO). For example, suppose a property has a cap rate (NOI/Value) of 8 percent and a lender makes a 75 percent LTV loan on the property. The loan yield would be $88\%/.75 = 10.67\%$. This suggests that if the lender took over the property in foreclosure, then it could recoup the amount owed to it on the loan provided it could sell the property at a cap rate of 10.67 percent or lower. (Keep in mind that price is inversely related to cap rate.) This gives the lender an indication that if prevailing market cap rates rose from 8 percent to 10.67 percent, the loan would still be recoverable (assuming the property’s NOI held up, and ignoring costs of foreclosure and sales). The loan yield doesn’t really provide any more information than the LTV, as it merely suggests that the property market could take a tumble in terms of its pricing per dollar of NOI equivalent to 25 percent (as $1/0.1067$ is 75% of $1/0.08$) and still cover the loan, which is exactly what is implied by the LTV of 75 percent. But people dealing with the property market often tend to think in terms of cap rates as the metric to track asset pricing, so it is an intuitive measure in that regard (even though it may be a heroic assumption to presume that the property NOI would remain the same in such circumstances).

Multiyear Proforma Projection. Although lenders typically scrutinize the ratio criteria described earlier (LTV, DCR, BER, and loan yield) most carefully for the initial year of the loan, they generally should require and examine a multiyear cash flow projection for the property and debt service,

“VC” as the variable costs (such that the total operating expenses can be expressed as the sum: $OE = FC + VC$), and assuming that VC is directly proportional to the occupancy in the building, then we have:

$$BER = (FC + DS) / (PGI - VC)$$

where VC is defined based on full occupancy (that is, $VC \cdot Occ$ is the actual variable cost component of the OE given an occupancy rate of Occ where Occ varies from 1.00 at full occupancy to zero at 100 percent vacancy). If the property is also subject to a ground lease (i.e., the property holding is actually a leasehold not a fee-simple ownership) and the ground lease requires regular annual payments, then add the amount “GL” to the numerator in the BER formula, where “GL” is the amount of the ground lease rental payment.

²⁵ In some cases, lenders attempt to define the NOI used in the DCR ratio to be net of recurring capital expenditures, such as leasing costs (commissions and TIs).

extending to the maturity of the loan. Borrowers usually supply such proforma projections, but lenders should take these with a grain of salt and attempt to double-check the assumptions. Lenders may examine any or all of the criteria noted previously for all the years in the proforma, not just the initial year.²⁶

The above list summarizes the major commercial loan underwriting metrics, the use of which in traditional underwriting is meant to be conservative and is often how they are in fact employed. As noted, even during the relatively rare periods when there is strong market pressure to “bend the rules,” that is, to make loans that are riskier than they should be (or than on paper they claim to be), the same metrics are still usually employed formally. During such times, the incentives for aggressive underwriting may be greater among originators who intend not to keep the mortgages but to “flip” them to a conduit pool in the secondary market. (It is widely believed that this occurred during the 2006–2007 bubble.) Securitization of mortgage pools sets up layers of structure between the ultimate investors in the bonds and the underlying individual borrowers in the loans, which makes it more difficult for originators who are not intending to keep “skin in the game” to resist pressures to be overly aggressive. But such market-wide systemic pressures can sometimes then reverse, after a crash, when lenders may become excessively conservative.

21.2.2 VARIABLES AND LOAN TERMS TO NEGOTIATE

Commercial mortgages present a potentially large array of possible loan terms and variables that can be negotiated. A partial list would include the following:

- Loan amount
- Loan term (maturity)
- Contract interest rate
- Amortization rate (in the extreme up to interest-only loans)
- Up-front fees and points
- Prepayment option and back-end penalties
- Recourse vs. nonrecourse debt
- Collateral (e.g., cross-collateralization)
- Lender participation in property equity

All such items can affect the risk and expected return in the loan. The number of these variables presents considerable flexibility and scope for creativity in the loan negotiation process. Typically, one or more of these items will matter more to one side of the deal than to the other side. As with any negotiation, much of the art of putting together a successful deal lies in finding a combination that maximizes both sides’ preferences. This requires an understanding of the trade-offs—most fundamentally, how changing the loan terms and variables affects the ex-ante risk and return on the loan from the perspective of both the borrower and lender. On the basis of previous discussions in this chapter and Chapters 19 and 20, you should have a good idea about the nature of this trade-off for some of the items listed. For example, we saw in Chapter 20 how reducing the amortization rate in the loan (i.e., longer amortization) can reduce the annual debt service (thereby increasing the DCR ratio) without changing the contract interest rate. Ultimately, the giving and taking in a successful negotiation result in ex-ante risk and return expectations in the loan that conform to the market and to the desires of both sides of the deal.

²⁶ See Chapters 11 and 14 for more detailed descriptions and discussion of the proforma.

21.2.3 NUMERICAL EXAMPLE OF THE NORMATIVE UNDERWRITING PROCESS

Our favorite real estate investor Bob wants to buy an existing fully operational 100,000-SF single-tenant office building in Iowa, and he has come to your firm, Sioux City Capital (Sioux for short), requesting a \$9,167,000 purchase-money mortgage. Bob's mortgage broker has put together a package indicating that Bob will be paying \$12,222,000 for the property, so that the requested loan would have an ILTV of 75 percent. Bob wants a 10-year interest-only nonrecourse loan. He is willing to accept a lockout loan that does not permit prepayment. You must decide whether to grant this loan as requested, reject it out of hand, or try to negotiate a modified loan.

To begin to answer this question, you first look at the capital markets to see what sort of contractual interest rate you would have to charge for a typical loan of this nature. You note that 10-year U.S. Treasury bonds are currently yielding 6 percent. You also note that in the commercial mortgage market spreads on nonrecourse 10-year lockout loans with a 75 percent ILTV are currently running at 200 basis points CEY. From this information, you can compute that the mortgage would have to carry a 7.87 percent interest rate (MEY).²⁷

Next, you consider the underwriting criteria Sioux is currently employing for loans of the type Bob wants. Given the fact that Sioux City Capital is actually a mortgage bank rather than a portfolio lender, Sioux will not be investing in Bob's mortgage on a long-term basis themselves. Rather, they will be placing into Bob's mortgage the capital of a large international life insurance company with whom Sioux has an ongoing relationship. It is this ultimate capital provider (who probably does intend to hold Bob's mortgage for the long term) whose underwriting criteria must be met. Unlike Sioux, this life insurance company does not know anything about Bob or the local real estate market, but they do know what sort of risk and return they are looking for in their commercial mortgage portfolio. Based on these considerations, they have specified the following underwriting criteria that Sioux will have to meet in any loans they issue:

1. Maximum ILTV < 75%
2. Maximum projected terminal LTV < 65%
3. In computing LTV: (a) apply direct capitalization with initial cap rate no less than 9 percent, terminal cap rate no less than 10 percent; (b) apply multiyear DCF valuation also for ILTV with discount rate no less than 10 percent; and (c) use the lower of the two ILTVs computed.
4. Minimum DCR > 120%
5. Maximum BER < 85% or average market occupancy less 5%, whichever is lower
6. Consider the need for capital improvements and avoid negative EBTCF projections.

With these requirements in mind, you examine Bob's loan request in some detail, applying your own knowledge of the space market in which Bob's property is situated. Bob's mortgage broker has submitted a package of information with the loan application including the following information about the property:

1. 100,000-SF fully leased single-tenant office building
2. Good-credit tenant (not a publicly traded corporation) signed a 10-year net lease three years ago.

²⁷ Note that spreads are quoted on Treasury bonds of equivalent maturity to the mortgage loan. Since mortgage and bond market yields must be equated on an effective annual rate (EAR) basis. Thus, the 200-basis-point spread implies that the mortgage must have a bond-equivalent yield of 8.00% (the 6.00% T-bond yield plus the 200-basis-point CEY spread). This equates to an EAR of 8.16%, which is therefore the loan interest rate that equates to the current capital market requirements. The monthly payments in the mortgage require a mortgage-equivalent yield of 7.87%, computed as $[(1 + 0.08/2)^{2(1/12)} - 1] \times 12 = 0.0787$. This is the rate we will use in the present illustration.

3. The lease has a current net rent of \$11/SF per year (with annual payments made at the end of each year), with step-up provisions to \$11.50 in year 5 (two years from now) and \$12.00 in year 8 (five years from now).
4. Current rents in the relevant space market are \$12/SF for new 10-year leases (with no concessions) and are expected to grow at a rate of 3 percent per year (e.g., will be \$12.36 one year from now, \$16.13 in 10 years).

The broker has also submitted the cash flow proforma projection shown in Exhibit 21-8. In this proforma, the broker has assumed a 75 percent probability of renewal for the tenant at the beginning of year 8 and has assumed that if the tenant does not renew there will be a three-month vacancy period. The broker has not included any provision for capital improvement expenditures and has assumed a 9 percent terminal cap rate to estimate the reversion value of the property at the end of year 10.

After examining this proforma and making some inquiries, you develop an alternative proforma projection for Bob's property that you feel is more realistic. You agree that the current market rent is \$12/SF, but you decide to reduce the expected rent growth rate assumption from 3 percent to 1 percent per year for Bob's building. Even though rents for new buildings in the space market might grow at 3 percent, Bob's building will be aging and becoming more obsolete in that market, so you feel that 1 percent is a more realistic expectation. After some investigation, you decide that the tenant, though not a publicly traded corporation, is a stable company that presents good credit risk. You decide to accept the broker's assumptions about this tenant's renewal probability and the vacancy downtime in the event of non-renewal, but you believe that some leasing expenses will be involved even if the tenant does renew. You decide to project year-8 leasing commission fees of \$2/SF if the tenant renews and \$5 if they do not renew. You also decide that it would be realistic to project tenant improvement expenditures of \$10/SF if the existing tenant renews and \$20/SF if a new tenant moves in. Finally, you decide that a 10 percent terminal cap rate would be a more realistic projection than the broker's 9 percent assumption considering that the building will be 10 years older by then.

Your modified proforma is shown in Exhibit 21-9, including the implied underwriting income ratios, the DCR and BER, given the \$9,167,000 interest-only loan that Bob wants, assuming the 7.87 percent interest rate currently required by the mortgage market.²⁸ It appears from Exhibit 21-9 that Bob's loan request will satisfy Sioux's income-based underwriting criteria. The initial DCR is 152 percent, well above the 120 percent minimum requirement, and the projected DCR is even higher in future years as the rent in the existing lease steps up and the space turns over at a higher projected market rent.²⁹ Your familiarity with the space market makes it clear that the initial BER of 60 percent is well below the average space market occupancy rate less 5 percent.³⁰

Unfortunately, Bob's loan proposal as it stands has several problems. One is apparent in the cash flow projection. It appears that there could well be a sharply negative EBTGF in year 8, the year the existing lease expires. The EBTGF is projected to be negative by over \$1 million in that year.³¹ This

²⁸ The monthly debt service on the interest-only loan is calculated as $(0.0787/12) \times \$9,167,000 = \$60,120$. Thus, the annual debt service is $12 \times \$60,120 = \$721,443$. (As with many numerical examples in this book, slight numerical discrepancies in equations are due to round-off.)

²⁹ The initial DCR of 152% is calculated as \$1.1 million NOI for year 1 divided by the \$721,443 debt service.

³⁰ The BER in Exhibit 21-9 has been calculated based on the projected market rent, rather than the property's rent based on the existing lease. This makes sense as the underwriting criterion compares the BER to the average occupancy rate in the market. Thus, the 60% BER for year 1 in the proforma is found as the \$7.21443/SF debt service divided by the \$12.12/SF projected market rent for that year.

³¹ The projected cash flow in year 8 reflects the mean or expectation across the renewal scenarios. It is calculated as follows. The vacancy allowance per SF is $(1 - 75\%) \times (0.25 \times \$12.99) = 0.81 = \$SF$. This is the nonrenewal probability ($1 - 75\% = 25\%$) times the conditional loss from vacancy that will occur if the tenant does not renew. This is expected to be three months of vacancy, or 25 percent of the income from a year that could otherwise earn the projected market rent of \$12.99/SF. The expected lease commission is $(75\%) (\$2/SF) + (1 - 75\%) - (\$5/SF) = \$2.75 = \SF . The expected tenant improvement expenditure is $(75\%) (\$10/SF) + (1 - 75\%) - (\$20/SF) = \$12.50 = \SF .

EXHIBIT 21-8

Broker's Submitted Proforma for Bob's Office Building

Year	1	2	3	4	5	6	7	8	9	10	11
Market rent (net)/SF	\$12.36	\$12.73	\$13.11	\$13.51	\$13.91	\$14.33	\$14.76	\$15.20	\$15.66	\$16.13	\$16.61
Property rent (net)	\$11	\$11.50	\$11.50	\$11.50	\$12.00	\$12.00	\$12.00	\$15.20	\$15.20	\$15.20	\$15.20
Vacancy allowance	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.95	\$0.00	\$0.00	\$0.00
NOI/SF	\$11.00	\$11.50	\$11.50	\$11.50	\$12.00	\$12.00	\$12.00	\$14.25	\$15.20	\$15.20	\$15.20
NOI	\$1,100,000	\$1,150,000	\$1,150,000	\$1,150,000	\$1,200,000	\$1,200,000	\$1,200,000	\$1,425,116	\$1,520,124	\$1,520,124	\$1,520,124
Reversion @ 9% cap											\$16,890,268

EXHIBIT 21-9

Sioux's Modified Proforma for Bob's Office Building and Loan Application

Year	1	2	3	4	5	6	7	8	9	10	11
Market rent (net)/SF	\$12.12	\$12.24	\$12.36	\$12.49	\$12.61	\$12.74	\$12.87	\$12.99	\$13.12	\$13.26	\$13.39
Property rent (net)	\$11.00	\$11.50	\$11.50	\$11.50	\$12.00	\$12.00	\$12.00	\$12.99	\$12.99	\$12.99	\$12.99
Vacancy allowance	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.81	\$0.00	\$0.00	\$0.00
NOI/SF	\$11.00	\$11.50	\$11.50	\$11.50	\$12.00	\$12.00	\$12.00	\$12.18	\$12.99	\$12.99	\$12.99
NOI	\$1,100,000	\$1,150,000	\$1,150,000	\$1,150,000	\$1,200,000	\$1,200,000	\$1,200,000	\$1,218,214	\$1,299,428	\$1,299,428	\$1,299,428
Lease commission	\$0	\$0	\$0	\$0	\$0	\$0	\$0	(\$275,000)	\$0	\$0	
Tenant improvements	\$0	\$0	\$0	\$0	\$0	\$0	\$0	-1,250,000	\$0	\$0	
Reversion @10% cap											\$12,994,280
Less OLB											\$9,167,000
PBTCF	\$1,100,000	51,150,000	\$1,150,000	\$1,150,000	\$1,200,000	\$1,200,000	\$1,200,000	(\$306,786)	\$1,299,428	\$14,293,709	
Debt service	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$721,443)	(\$9,888,443)
EBTCF	\$378,557	\$428,557	\$428,557	\$428,557	\$428,557	\$478,557	\$478,557	(\$1,028,229)	\$577,985	\$4,405,266	
DCR	152%	159%	159%	159%	166%	166%	166%	169%	180%	180%	
BER @ market	60%	59%	58%	58%	57%	57%	56%	56%	55%	54%	

problem can probably be dealt with, however. Single-tenant buildings, or properties in which a large proportion of the space is under leases that all expire near the same time, commonly face this type of occasional negative cash flow. Much of the problem is due to the need for capital improvement expenditures. The borrower may be able to finance some of this need. More to the point, there is ample positive EBTCF projected prior to the projected negative year, and the negative year is in the rather distant future by which time property appreciation should have reduced the danger of default. If necessary, some sort of sinking fund covenant might be included in Bob's mortgage to ensure that sufficient cash will be available to cover the projected shortfall in year 8.

A more serious problem in Bob's application is not apparent from the income analysis alone but appears in the valuation analysis. If one applies a 9 percent going-in cap rate to the initial year's NOI of \$1.1 million, the implied property value is \$12,222,000, the same price Bob is apparently planning to pay for the property, and just sufficient to allow the proposed \$9,167,000 loan to meet the 75 percent ILTV criterion. However, after investigating the relevant property market, you decide that realistic expected returns (going-in IRRs) for this type of property are well approximated by the 10 percent discount rate in your underwriting criteria. In any case, you do not see how you can apply a discount rate less than that in performing a DCF valuation of the property. At a discount rate of 10 percent, the 10-year projected property level cash flows (PBTCF) shown in Exhibit 21-9 (including the projected reversion) give a present value of only \$11,557,000, which implies a 79 percent ILTV.³² This violates the ILTV underwriting criterion.

As Bob's loan proposal does not fail your underwriting criteria by very much, you decide not to reject it out of hand, but rather to make a counterproposal. To meet the 75 percent ILTV limit in your underwriting criteria, you could offer Bob a slightly smaller loan of \$8.7 million. However, without amortization, such a loan would still not meet the **terminal LTV (TLTV)** criterion of 65 percent. Dividing \$8.7 million into the projected 10-year reversion value of \$12,994,000 results in a projected TLTV of 67 percent. However, the smaller loan could have some amortization and still produce an annual debt service payment similar to what Bob was originally asking for with his larger interest-only loan. For example, if the loan amortized at a 40-year rate, then the balloon payment after 10 years on an \$8.7 million loan would be \$8,230,047, which implies a TLTV ratio of only 63 percent, less than the 65 percent limit. The annual debt service on such a loan would be \$715,740 (with monthly payments), slightly less than the \$721,443 Bob was originally proposing.

This shows how trade-offs among the loan terms can be manipulated to meet underwriting criteria. Bob might accept a counterproposal of a smaller \$8.7 million loan with a 40-year amortization rate. If Sioux's underwriting criteria are competitive, and if your estimate is correct that the property market would not likely accept a realistic going-in IRR of less than 10 percent on a property like Bob's office building, then it should be difficult for Bob to find another lender that will lend any more than \$8.7 million on the property.³³

21.3 CHAPTER SUMMARY

This chapter introduced you to the basics of commercial mortgages, largely from the lender's perspective. You should now have a feeling for how to analyze the default risk in such loans and how underwriting is carried out in the commercial mortgage origination industry. In the next chapter we will step back a bit to consider the broader economics of commercial mortgages from the perspective of the ultimate investors in such loans.

³² $9,167,000 / \$11,557,000 = 79.3\%$.

³³ Admittedly, these are two rather large ifs. The difficulty of being able to observe or estimate property market values precisely, combined with lenders' desire to be competitive in the mortgage lending business, probably goes some way in explaining why the default rate is as high as it is, and why it has been so much higher on loans issued during boom times.

KEY TERMS

- Contract (stated) yield
- Expected return (ex ante yield)
- Coupon (coupon rate)
- Credit losses
- Yield degradation
- Recovery rate
- Loss severity
- Conditional yield degradation
- Hazard probability/function
- Conditional survival probability
- Cumulative survival probability
- Unconditional default probability
- Cumulative (lifetime) default probability
- Ex-ante (unconditional) yield degradation
- Ex-post yield degradation underwriting
- Initial loan-to-value ratio (ILTV) volatility
- Debt service coverage ratio (DCR)
- Break-even ratio (BER)
- Equity-before-tax cash flow (EBTCF) loan yield
- Terminal LTV (TLTV)

STUDY QUESTIONS

Conceptual Questions

- 21.1 What is the difference between the contract (or stated) yield and the realistic expected return (ex-ante yield) on a mortgage? Why is this difference important?
- 21.2 Ignoring multiple periods of time, what is the general relationship among the expected return, the contract yield (or YTM), the unconditional default probability, and the conditional yield degradation?
- 21.3 * What is the hazard function for a mortgage? What is the relationship between the hazard function, the unconditional default probability, and the cumulative default probability?
- 21.4 What is the major purpose of underwriting in the commercial mortgage industry? What is the relationship between underwriting and the market for commercial mortgage assets?
- 21.5 What are the two major foci of the lender's attention in commercial mortgage underwriting? Which one of these is usually more important, and why?
- 21.6 What is the difference between value-based and income-based underwriting criteria? In what way could you consider that the value-based criteria are more fundamental or important?
- 21.7 What is the relationship of the initial loan-to-value ratio (ILTV) to the default risk in the loan? How is this relationship affected by the volatility in the underlying property?
- 21.8 Why do lenders need to consider multiyear cash flow projections rather than just the initial income of the property?
- 21.9 What potential problem would the projected equity-before-tax cash flow (EBTCF) reveal that the DCR or BER would not reveal?

Quantitative Problems

21.10 In a one-period world, if the conditional yield degradation is 10%, the unconditional default probability is 15%, and the lender wants an expected return of 8%, what contract yield must the loan carry?

21.11 Consider a three-year mortgage with annual payments in arrears. Suppose the probability of default is 1% in the first year and 5% each year thereafter, given that default has not occurred previously.

- What is the hazard function of this loan?
- What is the unconditional default probability in year 2?
- What is the cumulative (or lifetime) default probability in this loan as of the time of its origination?

21.12 Suppose the loan in Question 21.11 is an 8% interest-only loan. If the conditional recovery rate is 75% each year, what is the expected return on the loan?

21.13 Suppose 10-year Treasury bond yields in the bond market are 7.00% CEY (or BEY), and the mortgage market requires a contract yield risk premium of 175 basis points (CEY). If a property has a net operating income (NOI) of \$400,000, and the underwriting criteria require a debt coverage ratio (DCR) of at least 125%, then what is the maximum loan that can be offered assuming a 30-year amortization rate and monthly payments on the mortgage?

21.14 Using the discounted cash flow (DCF) valuation method, what is the maximum loan that can be made on a property with the following annual net before-tax cash flow, assuming an 11.5% discount rate and underwriting criteria that specify a maximum loan/value ratio of 70%?

(Cash flows: \$1 million in year 1, 1.1 million in years 2 through 4, 1.5 million in years 5 through 9, and \$12 million in year 10 including reversion.)

21.15 A property has an expected first-year NOI of \$1 million. Recent sales of similar properties indicate that a first-year (or going-in) cap rate of 9.75% is reasonable for valuation purposes. A lender requires a minimum DCR (or DSCR) of 1.25 and will loan up to 70% of appraised value on a first mortgage. If the mortgage interest rate is 6.75%, payments are monthly, and the amortization period is 20 years, what is the maximum-sized loan the lender will advance?

21.16 The boxed feature “Is It Surprising that So Many Commercial Mortgages Default?” presents a simplified way to compute a rough estimate of the likelihood of mortgage default as a function of property value volatility and the LTV ratio of the loan when it is issued. There, it was argued that a 75% LTV ratio lending criterion would lead to approximately a one-sixth lifetime default probability. Use this same simplified approach, and the fact that about 5% of the normal probability distribution lies beyond two standard deviations from its mean, to estimate what LTV ratio lending criterion would be necessary to reduce the default probability to 1/40 (or 2.5%).

Part VI

Macro-Level Real Estate Investment Analysis

Most of this book up to now focused primarily at the micro-level analysis of individual properties and deals. While some of the preceding material included broad background information and fundamental building blocks relevant to the macro-level, the bulk of our attention on real estate equity investment so far has been at the micro-level.

That is about to change. Part VI will focus specifically on the major macro-level issues, concepts, and analytical tools. Here is where we study the forest as a whole rather than the trees. To be more precise, what we are referring to as the macro-level in real estate investment concerns the investor's decisions and management regarding many individual properties simultaneously, that is, aggregates of properties. The macro-level is also the level at which the investor's overall portfolio is considered, which is why it is often referred to as the portfolio level. At its broadest, this includes the investor's entire net wealth portfolio, including not just real estate but other investment asset classes as well, the so-called "mixed-asset portfolio." Indeed, it is at the macro-level that the interface between real estate and other asset classes enters most directly into analysis and decision making.

The concept of a macro-level of real estate investment analysis and management really only dates from the last third of the twentieth century. Traditional real estate investment was effectively a purely micro-level endeavor. Indeed, the distinction between these two levels is somewhat unique to real estate. The concepts of macro-level investment decisions and activities, and macro-level valuation, for example, are a bit foreign to the fields of securities investments and corporate finance. Yet during the last two decades of the twentieth century, the macro-level of real estate investment analysis and decision making blossomed rapidly into a major component and force within real estate, with its effects permeating down to even small-scale microlevel transactions, not least because the macro-level is central to the growing link between Wall Street and Main Street. The macro-level is therefore also the level where we find many of the more modern and rapidly growing real estate professional career paths.

At a broad-brush level, and by way of an initial introduction, it is useful to think of three major types of macro-level real estate investment decision arenas: strategic policy formulation, tactical policy formulation, and policy implementation. Strategic investment policies define broad, overall allocations and long-run directions and objectives. Tactical policies seek to profit from shorter-term

opportunities. Implementation concerns how to carry out policies of both the strategic and tactical type most effectively. For example, an analytical tool widely used at the strategic level is modern portfolio theory (MPT). A major concern at the tactical level is market timing and econometric forecasting of the space and asset markets. Major topics in implementation include the analysis and evaluation of the performance of investment managers and the crafting of incentive structures that align the interests of such managers with their investor clients. In all three of these decision areas, the macro-level cannot be divorced from the micro-level. Macro-level decisions are ultimately implemented at the micro-level, and the quality of the macro-micro link is a key to long-run success in real estate investment.

We will address all three of these macro-level decision arenas in Part VI, as well as the macro-micro link, by presenting some fundamental macro-level principles and analytical tools and showing how these relate to investment decision making. The major macro-level investment principles and analytical tools we will address include portfolio theory, equilibrium asset pricing theory, macro-level valuation and return measurement issues, and real estate investment management performance attribution and evaluation.

Part VI is organized into four chapters. Chapters 22 and 23 begin with a presentation of MPT and equilibrium asset pricing models, the classical tools and discipline for “top-down” real estate investment decision making, including particular consideration of the role of real estate in the mixed-asset portfolio. In Chapter 24, we get down to a very fundamental nuts-and-bolts problem, that of dealing with the peculiarities and measurement issues inherent in real estate investment performance data, including some cutting-edge tools for how to construct real estate price indexes, a key and fundamental exercise for modern real estate investment analysis. Finally, Chapter 25 focuses on the investment policy implementation decision arena with a detailed look at the real estate investment management industry, including quantitative performance attribution and evaluation.

22 Portfolio Theory

Strategic Investment Considerations

CHAPTER OUTLINE

- 22.1 Intro to Mean-Variance Portfolio Theory, MPT
 - 22.1.1 Investor Preferences and Dominant Portfolios
 - 22.1.2 Portfolio Theory, Diversification, and the Role of Real Estate in the Portfolio
 - 22.1.3 The Efficient Frontier
 - 22.1.4 Bringing in Investor Preferences
 - 22.1.5 Major Implications of Portfolio Theory for Real Estate Investment
- 22.2 Allowing for a Riskless Asset: The Two-Fund Theorem and the Sharpe-Maximizing Portfolio
 - 22.2.1 Two-Fund Theorem
 - 22.2.2 Sharpe-Maximizing Portfolio
 - 22.2.3 Summary of the Implications of the Riskless Asset Construct
- 22.3 Other Models of Portfolio Allocation: Risk Parity and Passive Allocation
- 22.4 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- What is meant by modern portfolio theory (MPT) and how to apply this theory using computer spreadsheets.
- The major strategic investment policy implications MPT holds for real estate at the broad-brush level of the overall mixed-asset portfolio.
- The usefulness of the riskless asset assumption and the meaning of the Sharpe-maximizing portfolio.
- The difference between Mean-Variance Portfolio Theory (MPT) and other approaches to portfolio allocation widely used among professional investors: Passive allocation, and Risk parity.

Throughout the financial economics part of this book, we emphasized two major concerns of investors: total *return* and the *risk* surrounding that return. In this chapter, we take our decision-making focus to the level of the investor's overall wealth, the portfolio level. At that level, the most basic investment decision is how to allocate the investor's wealth, how much should be invested in one asset, or one *class* of assets, versus another, among the range of all the possibilities. This is a so-called "top-down" investment decision because it provides the high-level strategic direction for the investments, rather than the specific decision of which individual investment asset to invest in or how much to pay for any given investment ("bottom-up" decisions). But return and risk still remain the major foci of our attention. These are the investor's main concerns, whether at the "macro" level, where we are here in Part VI or at the "micro" (individual property) level, where much of the earlier part of this book has focused.

In practice for real estate investing, such high-level strategic investment decision making is based on formal analysis and methodology typically only used by high net worth family offices and large investment institutions such as pension and endowment funds, sovereign wealth funds, life insurance companies and so forth. Broadly speaking, three major formal methodologies are prevalent: index or benchmark based passive allocation, risk parity allocation, and mean-variance portfolio theory (aka “Modern Portfolio Theory” – MPT). Of these three, MPT has the most rigorous analytical grounding and academic pedigree. It is traditionally the most widely taught in graduate schools of business, and it is widely considered in the real world. Accordingly, we will devote most of this chapter to educating you in this Nobel Prize-winning theory and toolkit. But the other two approaches are also very widely employed and perhaps ultimately are closer to determining the actual allocation strategies of many institutions. Passive allocation and risk parity can be presented more easily than MPT as they have less deep and rigorous bases, and we will leave our briefer treatment of them to the end of the chapter. We turn first to MPT.

22.1 INTRO TO MEAN-VARIANCE PORTFOLIO THEORY, MPT

To address the wealth allocation decision in a rigorous way, MPT makes three major contributions: (1) it treats risk and return together in a comprehensive and integrated manner; (2) it quantifies the investment-decision-relevant implications of risk and return; and (3) it makes both of these contributions at the portfolio level, the level of the investor’s overall wealth.

MPT’s elegance allows it to reveal, simply, some fundamental aspects of the world, aspects that cut right to the heart of good strategic investment decision making. Learning this theory can build your fundamental investment expertise and sophistication. And this is why, from its inception, MPT has been used directly and widely in the real world of investment practice. This is not just ivory-tower academic stuff. It provides insights that are important for investment decision-making at the broadest and most fundamental level. These insights are especially important for real estate, as we will see.

Modern Portfolio Theory is also often referred to as **Mean-Variance Portfolio Theory**, or as **Markowitz Portfolio Theory**, for the economist who is largely credited with its development (and won a Nobel Prize for it). Conveniently, all three of these names can be abbreviated MPT. In MPT the objective of the macro-level investment decision maker is taken to be the minimization of portfolio volatility (or variance) subject to an expected (mean) return target, or equivalently, the maximization of portfolio expected return subject to a volatility constraint. The return (which refers to *total* return, income plus capital growth, on the investor’s total wealth) is clearly what the investor wants. And volatility (the standard deviation over time in the return) is taken to measure the risk that the investor is concerned about. Recall from Chapter 3 that total return is the combination of returns from cash flows and net appreciation expected for a given holding period, and risk (volatility) is typically measured by historical data, or using forward looking indicators when available.¹ To grasp the rigorous nature of MPT and begin to build your intuition, let us step back and look at the basis of the theory.

22.1.1 INVESTOR PREFERENCES AND DOMINANT PORTFOLIOS

As risk and return are the two main issues MPT is concerned about, a good place to begin is to consider a picture of the risk and return preferences of the investor. Exhibit 22-1 depicts risk and return on the horizontal and vertical axes of a rectangle, respectively. Each point in the rectangle is a different combination of risk and return. The investor’s preferences are indicated by the indifference curves shown in the rectangle. These curves show how the investor judges the trade-off between risk and return for her investments and her degree of risk aversion. The investor is indifferent among portfolios that provide risk and return combinations that lie on the same indifference curve.

¹ Forward-looking volatility measures require the use of futures such as options on future contracts, few of which exist in the real estate market.

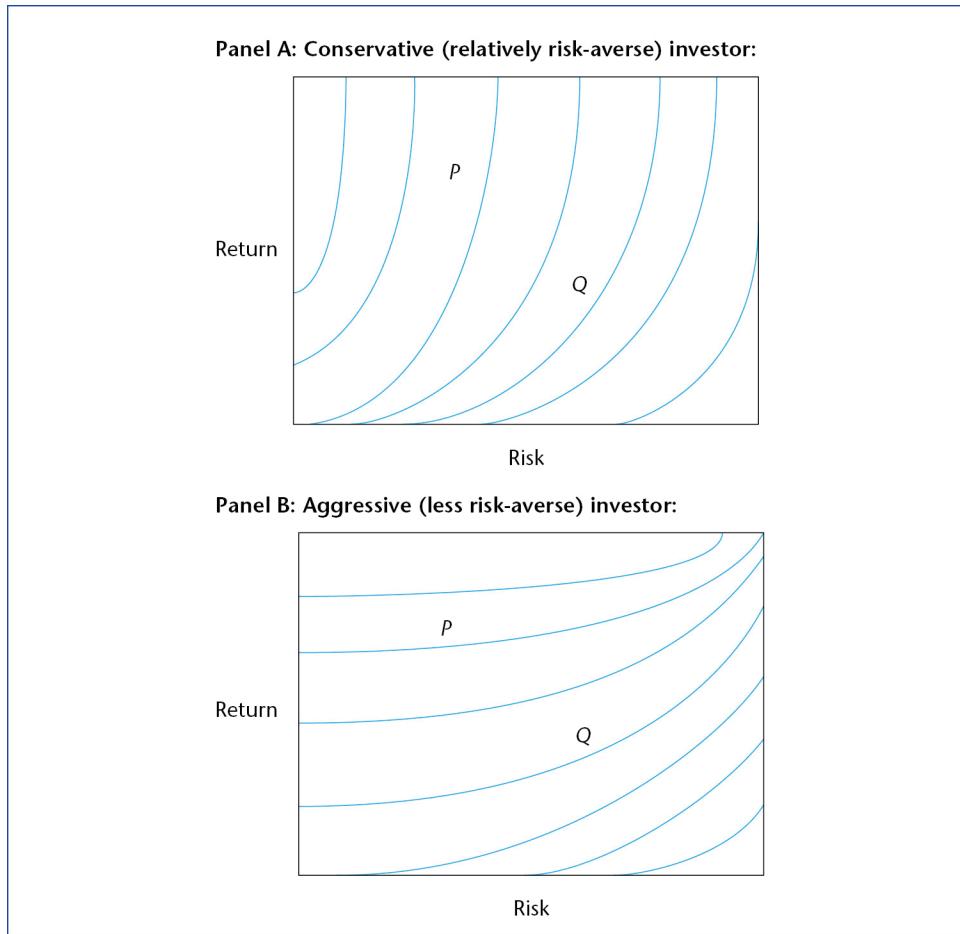


EXHIBIT 22-1 Utility Preference Surface (indifference curves) of an Investor.

You can think of the indifference curves as contour lines on a map of a ski slope, representing a three-dimensional surface above the risk/return rectangle. Higher points on this surface indicate risk/return combinations that are preferred by the investor. Other things being equal, investors generally prefer a greater return and less risk. Thus, the investor preference surface is rising toward the upper left corner of the rectangle. The investor will prefer points farther to the “north” and “west” in the rectangle; points such as *P* will be preferred over points like *Q*.

The rectangle in the top panel in Exhibit 22-1 (panel A) depicts a conservative investor, one who is relatively more risk averse (less risk tolerant) in his investment preferences. Such an investor’s indifference curves are steeply curved to the north in the risk/return rectangle, indicating that he must be compensated with a lot of additional return in order to be willing to take on a little more risk. This is seen by the fact that the curves move a considerable distance along the vertical (expected return) axis for each small increment along the horizontal (risk) axis.

In contrast, the rectangle in panel B of Exhibit 22-1 depicts a more aggressive investor, one who is relatively more risk tolerant. Such an investor’s indifference curves are less steeply rising over the horizontal axis, indicating less need for additional expected return in more risky investments.

Suppose that points *P* and *Q* in Exhibit 22-1 represent the expected return and risk of two different portfolios, that is, two different allocations for the investor’s wealth. Clearly, both the conservative and the aggressive investors depicted in Exhibit 22-1 would prefer portfolio *P* to portfolio *Q*. In fact,

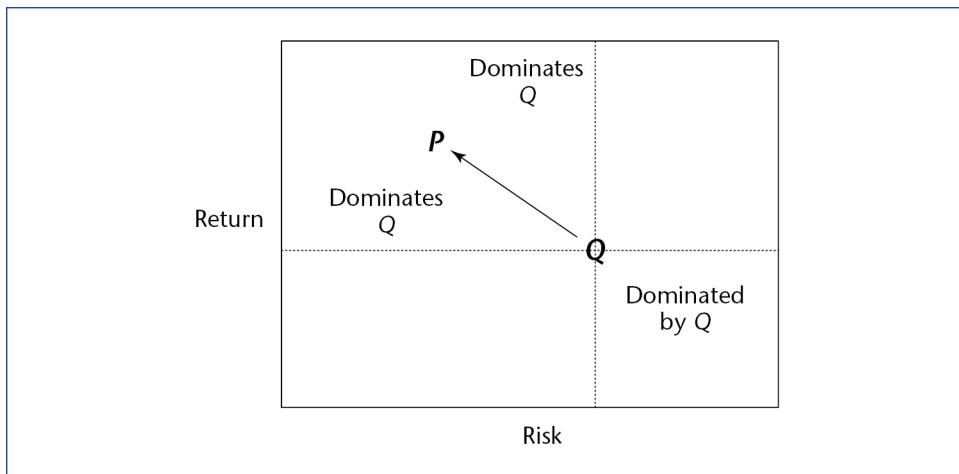


EXHIBIT 22-2 Portfolio Dominance.

as long as investors are not actually “risk loving” (that is, as long as they do not actually prefer more risk to less, holding return constant), they will always prefer a point such as P to one like Q .

Now consider Exhibit 22-2. Here, the risk and return possibilities are divided into four quadrants emanating from portfolio Q . Any portfolio (such as P) that provides as much or more expected return than Q with as little or less risk, is said to *dominate* Q . This would be any portfolio providing risk/return combinations above and/or to the left of Q , that is, in the upper-left quadrant. Similarly, any risk/return combination in the lower-right quadrant would be *dominated by* Q . Any investor, no matter what her **risk preferences** (whether she is conservative or aggressive, as long as she is not actually risk loving), will prefer a **dominant portfolio** to a dominated portfolio.

In practice, and from a quantifiable perspective, MPT is used primarily to help investors avoid holding dominated portfolios. MPT is about moving investors from points such as Q to points such as P whenever possible, that is, moving up and/or to the left (or “northwesterly”) in a standard risk/return diagram such as that in Exhibit 22-2. MPT has less to say (quantitatively) about choosing between Q and another portfolio that lies in either the quadrant to Q ’s northeast or the quadrant to Q ’s southwest. To rigorously make such a choice, we would need to know how to quantify the investor’s risk preferences more completely than simply knowing that he is not risk-loving.²

² Some techniques have been developed to attempt to quantify investor preferences more precisely. For example, the investor may be guided through an exercise in which she states her preference for various combinations of stocks and bonds given various different expected returns and historical volatility patterns, in such a way that a “preference map” along the lines of what is depicted in Exhibit 22-1 can be constructed for the investor regarding various risk/ return combinations. More commonly in the real world, however, this part of the portfolio allocation decision (that is, movement in a southwesterly or northeasterly direction in the risk/return diagram) is treated “heuristically” or non-quantitatively, which is to say for many investment institutions, politically. Implicitly, of course, some indication of the investor’s preferences is inevitably revealed in the allocation policy he finally adopts. That the adoption process should be somewhat *political* actually makes *economic* sense. Although a Nobel Prize was given to Harry Markowitz for his development of portfolio theory, another Nobel Prize was given to economist Kenneth Arrow, in part for his development of what is called the “impossibility theorem.” Arrow proved that it is mathematically impossible to define a so-called social welfare function that can be relied on to aggregate individual preferences into a single rational preference function. In effect, viewed from a very fundamental perspective, any investment decision-making institution that has more than one individual owner or beneficiary must employ some sort of political process to aggregate preferences effectively, so as to arrive at a common decision. Alas (but aren’t you really glad), there is a limit to how far technocrats can take us.

22.1.2 PORTFOLIO THEORY, DIVERSIFICATION, AND THE ROLE OF REAL ESTATE IN THE PORTFOLIO

The expression “Don’t put all your eggs in one basket” has been around a long time. Long before Harry Markowitz published his seminal paper on portfolio theory in 1952, people had an intuitive sense that they should not put too much of their total wealth into a single investment or type of asset. This might be referred to as intuitive diversification. It is common sense that if all (or nearly all) of your wealth is invested in only one type of asset, or one industry or sector, then you are overly exposed to loss in the event of a downside event that randomly affects only that one type of asset or industry.

Portfolio theory adds to this primitive concept of diversification by quantifying the benefit of diversification in terms of portfolio risk and return, and by providing some rigorous (albeit somewhat simplistic) guidance as to exactly how to diversify, that is, *how many eggs should you put in which baskets*.

In this regard, it is important to understand that some combinations of assets are more valuable than others as far as their diversification effect is concerned. For example, many individual stocks and sectors of the stock market are highly positively correlated over time.³ But in general, broad asset classes tend to be less positively correlated. For example, stocks and bonds tend to be less correlated than individual assets or groups of assets within each of those two asset classes or within the same industries. Real estate, in the form of direct private property investments (as distinguished from REITs) forms a separate asset class, distinct from either stocks or bonds. As such, real estate investments in general tend to be less positively correlated with stocks or bonds than will typically be the case among investments within any one such asset class.

The key mathematical point is that pairs of assets or asset classes that do not move together provide a greater diversification benefit when they are combined in a portfolio. The volatility of such assets tends to cancel out when they are combined in a portfolio. Two investments each of which has relatively low volatility (return variation across time) but which are highly positively correlated (tend to move together), might result in higher volatility when combined together in a portfolio than two investments that have relatively higher volatility each but with lower correlation between them. MPT allows us to apply this basic principle to modern databases of historical investment returns of different asset classes, or to forecasts of future returns, to analyze rigorously optimal portfolio allocations.⁴

There is a famous adage about economics that one of the things “the dismal science” is always propounding is that “*there is no free lunch!*” But diversification may be one of the rare exceptions that proves that rule. Diversification can be a type of “free lunch.” It allows the investor to “get something for nothing,” in the sense that it can allow you to improve your expected return without increasing your risk, or to reduce your risk without reducing your expected return.

The level of strategic investment planning dealing with an investor’s entire wealth, or with the entire portfolio that an institution is responsible for, is referred to as the “**mixed-asset**” perspective. The choice set is not limited only, for example, to stocks, or bonds. This is the level at which the greatest advantage can be taken of the benefits of diversification described above. This perspective is very important for real estate investment. As noted, after stocks and bonds,

³ The definition and meaning of the correlation coefficient is explained in Appendix 22. Appendix 22 is available online in the Chapter supplemental materials.

⁴ One might also intuitively understand that assets that are negatively correlated or in the extreme, perfectly inversely correlated, combine in such a way as to dramatically lower overall portfolio volatility when combined in the right proportions. Such asset return correlations are hard to find in the real world, but negative correlations do exist, in particular by the creation of derivative products created for the purpose of hedging. But, “*there is no free lunch*,” hedging reduces the expected return in the combined hedged portfolio.

real estate comprises the third major asset class, representing a large amount of physical capital with unique risk and return characteristics. As such, real estate gives investors the opportunity to diversify their portfolios more broadly than they could with only two major asset classes. This **diversification benefit** of real estate is a major reason for the substantial inclusion of real estate in many large institutional portfolios.

TEXT BOX 22.1 ANOTHER ROLE FOR REAL ESTATE

MPT highlights real estate's role as a diversifier within the portfolio. This role makes sense from the MPT perspective, viewing the mixed asset portfolio as encompassing the investor's entire net wealth. From the perspective of a pension fund, however, or other such institution managing a portfolio of assets at least in part to meet future liability obligations, a somewhat different perspective may be warranted. The net wealth of the investor may be viewed as the asset portfolio minus the present value of future liability obligations (e.g., pension benefit payments to retirees). These obligations are generally largely contractually fixed (in the case of a "defined benefit" pension plan, for example). Managers of the fund must therefore consider two objectives, not only the risk and return optimization of the fund's assets, but also its need to meet obligations on the liability side. Although bonds are particularly useful in such strategies, real estate also has a useful role on the liability side. In particular, pension benefit obligations tend to be sensitive to inflation because of **cost-of-living adjustments (COLAs)** in the benefits. As noted in previous chapters, real estate is a relatively good inflation hedge, whereas long-term bonds are notoriously exposed to inflation risk. Thus, if real estate's major role on the asset side of the pension fund's balance sheet is as a diversifier, its major role to help with the liability side is often as an inflation hedge asset.⁵

22.1.3 THE EFFICIENT FRONTIER

Let us continue our explication of MPT by putting together the general point about diversification we made in Section 22.1.2 with the idea described in 22.1.1 about avoiding dominated portfolios. This synthesis is one of the most important ideas in MPT: the efficient frontier.

To understand the concept of the efficient frontier, let's look at the big picture, a "30,000-foot perspective", by supposing there are only the three most basic investment asset classes: stocks, bonds, and real estate. First, consider the risk and return opportunities if we could only invest in one asset class, that is, if not diversify across asset classes, but instead if we had to choose one "pure" portfolio, either stocks, bonds, or real estate. And let us suppose for illustrative purposes that we face the risk/return expectations given in Exhibit 22-3. Then we could either get 6 percent expected return

⁵ That real estate assets provide a good inflation hedge is common sense based on the fact that real estate's value is based on a "real" physical product. (See also Section 15.2.4 in Chapter 15 for the mechanics of how this can interact with debt finance.) However, the empirical evidence for how well or how strongly real estate returns are positively correlated with inflation is mixed. Anecdotally, real estate performed well in nominal terms in the late 1970s and early 1980s when inflation was most problematic in U.S. history. Certain types of real estate (including single-family housing) also performed well in the post-pandemic inflation of the early 2020s.

EXHIBIT 22-3
Illustrative Risk and Return Expectations

	Stocks	Bonds	Real Estate
Expected Return	10.00%	6.00%	7.00%
Volatility	15.00%	8.00%	10.00%
Correlation with:			
Stocks	100.00%	30.00%	25.00%
Bonds		100.00%	15.00%
Real Estate			100.00%

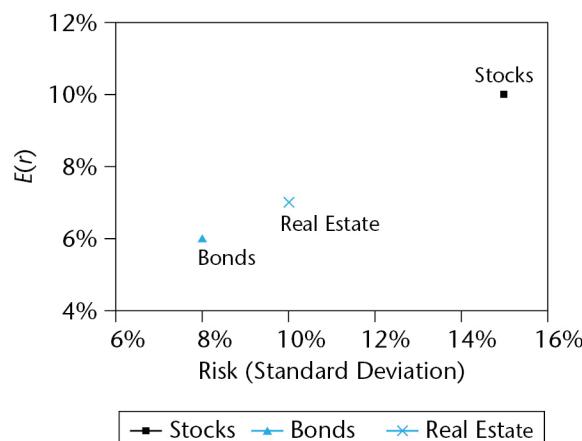


EXHIBIT 22-4A Three Assets: Stocks, Bonds, Real Estate (no diversification).

and 8 percent expected volatility by investing in bonds, or 7 percent return and 10 percent volatility in real estate, or 10 percent return and 15 percent volatility in stocks. These three possibilities are illustrated by the three points in Exhibit 22-4A.

Notice that these three possibilities all lie generally in a northeasterly/southwesterly relation to one another in the risk/return diagram. No one possibility is dominated by any other. This is typical of the risk/return relationships we would expect among asset classes in a well-functioning capital market. Asset classes that are dominated would be shunned by investors, driving down their prices and driving up their returns until they were no longer dominated. Similarly, investors would flock to a dominant asset class, driving up its prices and down its returns until it no longer dominated. But the possibilities shown in Exhibit 22-4A are highly constrained, not allowing any diversification across asset classes.

Suppose we relax this constraint a little by allowing diversification across any two asset classes. This expands the risk/return possibilities to the three curved lines indicated in Exhibit 22-4B. Each curve represents the risk/return possibilities from mixing two asset classes, either stocks and bonds, stocks and real estate, or bonds and real estate. For example, the curve on the lower left, connecting the pure bond and pure real estate possibilities, represents risk and return combinations available from various allocations to bonds and real estate.

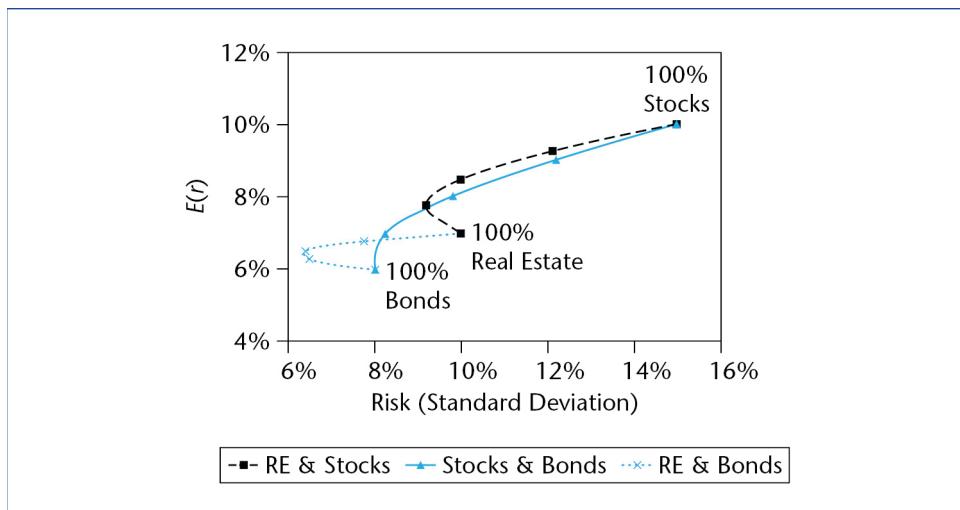


EXHIBIT 22-4B Three Assets: Stocks, Bonds, Real Estate (with pairwise combinations).

Now we see that certain pairwise combinations are dominated by others. Some points on each of the three curves lie to the northwest of some points on the other curves. For example, a portfolio containing 70 percent stocks and 30 percent real estate dominates a portfolio with 75 percent stocks and 25 percent bonds. The former has an expected return of 9.1 percent with volatility of 11.6 percent, while the latter has an expected return of 9.0 percent and volatility of 12.0 percent, that is, a lower expected return with more risk. An investor with the risk and return expectations presented in Exhibit 22-3 should never hold a portfolio of 75 percent stocks and 25 percent bonds if they could invest in real estate instead of bonds.

It is also interesting to notice the shape of the three curves in Exhibit 22-4B. They all curve up and to the left as compared to a hypothetical straight line connecting the two endpoints of the curve (pure single asset class investments). This means that diversified mixtures of any two asset classes will provide some additional expected return for a given amount of risk, or less risk for a given expected return, between the risk/return limits implied by the undiversified (pure) investments (the endpoints of the curves). This favorable “bending” of the risk/return possibility curve is a mathematical result between any two assets or asset classes that are not perfectly positively correlated. This is important because in reality, no two *underlying* assets (or asset classes) will ever be perfectly correlated. The result we see in Exhibit 22-4B is typical in this regard.

Note also that the curves involving real estate combinations (either with bonds or stocks) are more sharply bent to the left than the curve that represents stock and bond combinations. This is the effect of real estate’s relatively low correlation with the other two asset classes, especially with bonds. In fact, among the three pairs depicted in Exhibit 22-4B, the most favorable bending is in the real estate/bond combinations. Note that the curves in the Exhibits are based on the assumed correlations among the asset classes given in Exhibit 22-3: +15% between real estate and bonds, +25 percent between real estate and stocks, and +30% between stocks and bonds.

Exhibit 22-4B still represents a constraint on the investor’s ability to diversify, as diversification is limited to only two asset classes. Suppose we relax this restriction. With the possibility of diversifying across all three asset classes, many more risk/return possibilities open up. As no investor should want to hold any portfolios that are dominated, the set of interesting combinations of the three asset classes consists only of those that are not dominated. Any such portfolio is called an

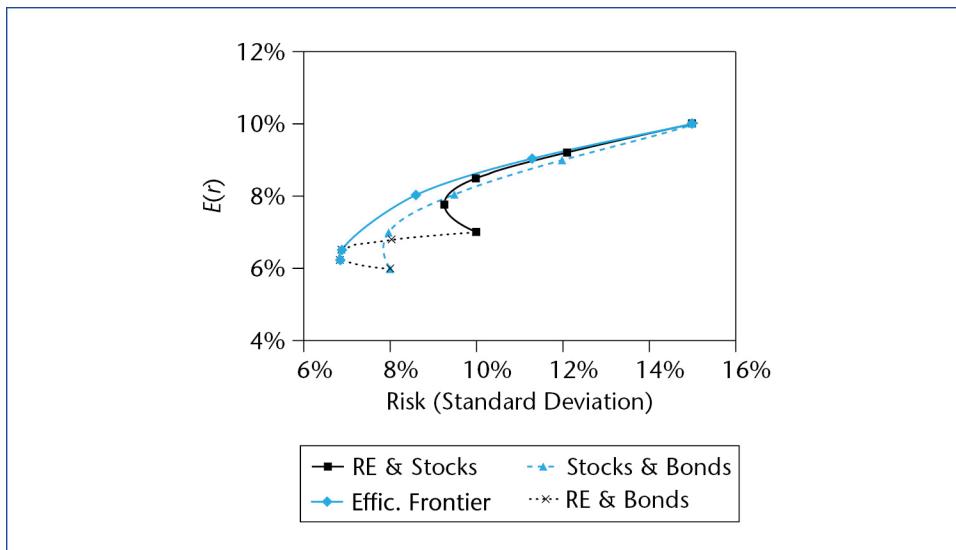


EXHIBIT 22-4C Three Assets with Diversification: The Efficient Frontier.

efficient portfolio, and the risk/return possibilities associated with the set of all possible efficient portfolios is called the **efficient frontier**.⁶ Any portfolio on the efficient frontier is, by definition, not dominated by any other portfolio, and any non-dominated portfolio is on the efficient frontier. In other words, the efficient frontier consists of all asset combinations that maximize return and minimize risk. In other words, for a given volatility, the efficient frontier maximizes the expected return, and for a given expected return the efficient frontier minimizes the volatility.

The efficient frontier for combinations among the three asset classes in our numerical example is shown in Exhibit 22-4C.⁷ The solid blue curved line marked with diamonds is the efficient frontier for stocks, bonds, and real estate given the risk and return expectations indicated in Exhibit 22-3. This curved line represents the best risk and return possibilities, those that are farthest northwest in the diagram. Points farther to the north or west of the efficient frontier are not feasible. Now you see why this is called a frontier. The efficient frontier always has this characteristic curved shape, running generally from the southwest toward the northeast in the risk/return diagram, bending upward and leftward relative to a straight line connecting the two endpoints.

Each point on the efficient frontier represents a unique combination of its possible constituent assets. In the case represented in Exhibit 22-4C, each point on the frontier represents a unique proportional allocation among stocks, bonds, and real estate.⁸ Each point on the efficient frontier also

⁶ **Terminology alert:** Do not confuse *efficient* here with the sense of the term used in financial economics to describe the functioning of an asset market, that is, the concept of *informational efficiency*. An efficient portfolio is simply one that is not dominated by any other portfolio in its expected return and risk. This does not carry any implication regarding the informational efficiency of the asset markets in which its constituent assets trade.

⁷ The efficient frontiers shown in Exhibits 22-4C and Exhibits 22-5 and 22-6 have been calculated using a simple Excel® template that is available in the online supplemental materials for this text. You can use this template to try out your own risk and return expectations and perform sensitivity analysis on our results presented here. The Chapter 22 supplemental materials include a screencast video that walks through how to use the Excel template.

⁸ In general, some of the efficient portfolios may have zero or negative allocations to one or more constituent assets. In the example we have been considering here, none of the allocations is negative because we have constrained our portfolio to avoid short positions (which gives the frontier a finite upper endpoint). The “no shorts” constraint is generally realistic for large institutional portfolios when considering allocation strategy across broad asset classes.

corresponds to a unique risk/return combination. For example, the point on the efficient frontier corresponding to an expected return of 7 percent entails a volatility of 6.89 percent and represents an allocation of 16 percent to stocks, 48 percent to bonds, and 36 percent to real estate.

22.1.4 BRINGING IN INVESTOR PREFERENCES

Remember our point in 22.1.1 about dominated and non-dominated portfolios. No investor should want to hold a dominated portfolio. Portfolio theory tells us that all investors should hold portfolios on the efficient frontier. Which portfolio on the efficient frontier any given investor should hold depends on the risk and return preferences of the investor. As noted, the efficient frontier is generally sloped from the southwest to the northeast in the risk/return diagram. A rational choice among different points on the efficient frontier can only be made by consulting the investor's preferences for risk and return. In practice these preferences are often expressed in terms of a target long-run rate of return for the portfolio, and within an institution this target may be selected by a more or less political process, often reflecting the policy needs and constraints of the institution.⁹ In effect, investors who are more aggressive or risk tolerant specify higher return targets (knowing that the capital market will dictate that this will imply greater portfolio volatility), while conservative investors specify lower return targets.

This process is depicted in Exhibit 22-5. Both panels show the efficient frontier we just constructed based on the risk and return expectations in Exhibit 22-3. The other curve shown in the exhibit is an

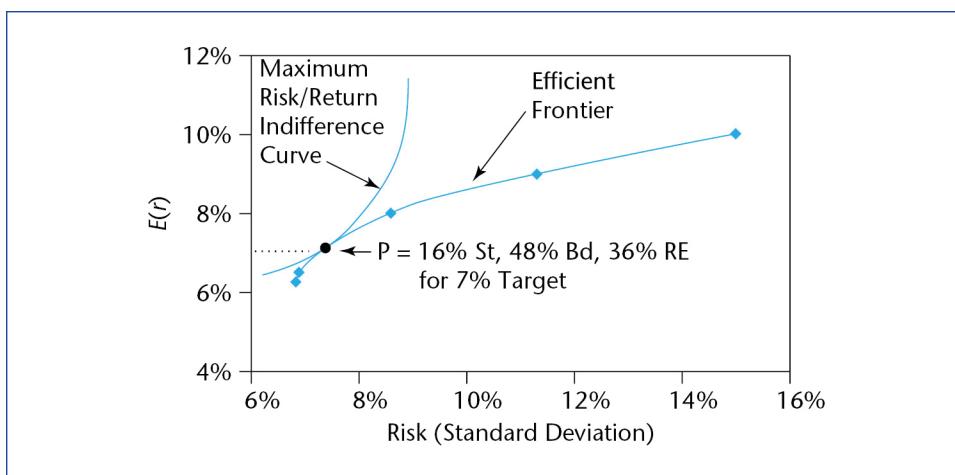


EXHIBIT 22-5A Optimal Portfolio for a Conservative Investor ($P = 7\%$ target).

⁹ A relatively sophisticated way to help make this decision which has become popular is called "Value at Risk" (VaR). In a VaR analysis a Monte Carlo simulation model is developed to simulate how the portfolio might perform over time (based on assumed expected returns and covariances specified for all of the candidate asset classes). Various different portfolio allocations are tested via the simulation model with a focus on what probability, based on the simulation results, is associated with a loss in portfolio value of a given specified magnitude. Of course, the VaR analysis is just as susceptible to the GIGO ("garbage in, garbage out") weakness as any analytical model. It may make the future appear less uncertain or more understood than it really is. A particular challenge is to capture the complete shapes of the future return distributions and covariances, including possible "fat tails" (nonnormal probability distributions that have greater probability of extreme outcomes than implied by normal probability) and covariances that increase to abnormal levels during financial crises. These types of problems are swept under the rug in classical MPT, which assumes stable and finite covariances. But at least the classical model may be less susceptible to lulling decision makers into a false impression of greater knowledge than they actually have.

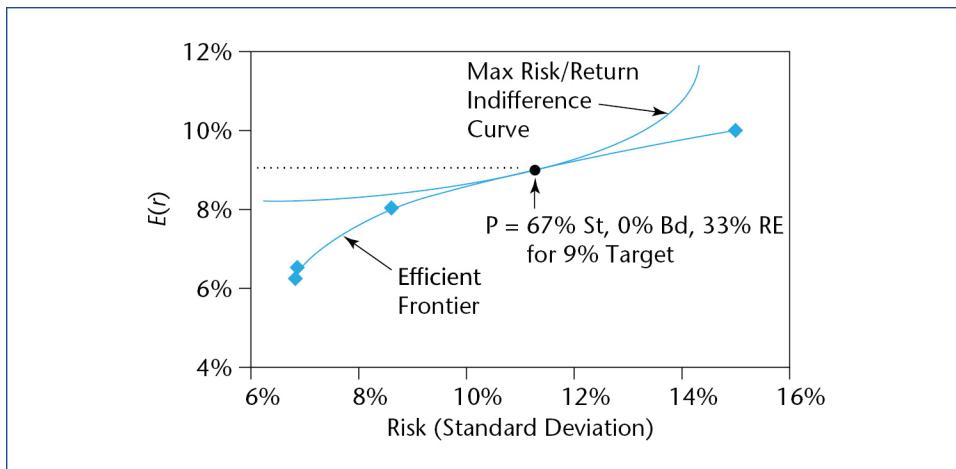


EXHIBIT 22-5B Optimal Portfolio for an Aggressive Investor ($P = 9\%$ target).

indifference curve reflecting the investor's preferences, similar to the isoquants in Exhibit 22-1. In particular, the indifference curve reflects the highest level of satisfaction (or "utility") the investor can achieve. Indifference curves that lie parallel to the curve shown would either be infeasible (to the northwest of the depicted curve) or indicative of lower levels of investor satisfaction (to the southeast of the depicted curve).¹⁰ The indifference curve that is just "tangent" to the efficient frontier (i.e., the indifference curve that touches the efficient frontier at one and only one point) is the best the investor can do. In this sense, the investor's selected target return (on the vertical axis, or equivalently their corresponding target volatility on the horizontal axis) reflects the investor's risk/return preferences, as described in Section 22.1.1.

The top panel, 22-5A, shows this result for a relatively conservative investor, indicating a target return of 7 percent. The bottom panel, Exhibit 22-5B, shows the result for a more aggressive investor, indicating a target return of 9 percent. The 9 percent target implies an optimal portfolio volatility of 11.30 percent as compared to the 6.89 percent volatility implied by the 7 percent target. While the efficient portfolio for the 7 percent target had only 16 percent allocation to stocks, the efficient portfolio for the 9 percent target would have 67 percent allocation to stocks, based on the same risk and return expectations (those of Exhibit 22-3). The 9 percent target has a 33 percent allocation to real estate, and zero percent to bonds.

22.1.5 MAJOR IMPLICATIONS OF PORTFOLIO THEORY FOR REAL ESTATE INVESTMENT

Exhibit 22-6 shows another way to present the efficient frontier. This Exhibit is what is called an "area chart," which shows how the optimal shares of the constituent asset classes change as a function of the investor's expected return target (i.e., as a function of the investors' risk preferences). The portfolio target rate of return is indicated now on the horizontal axis (instead of the vertical axis as before). The vertical dimension of the rectangle in Exhibit 22-6 stacks the efficient allocations to each asset class on top of each other to account for 100 percent of the investor's wealth. A vertical line is like a transect, indicating the share of each asset class in

¹⁰ Recall that an indifference curve (also called an "isoquant") is like a contour line on a topographic map. No two indifference curves (representing different levels of utility of the same investor's preferences) can ever cross each other. So indifference curves mapping out an investor's preferences are parallel to each other, as in Exhibit 22-1.

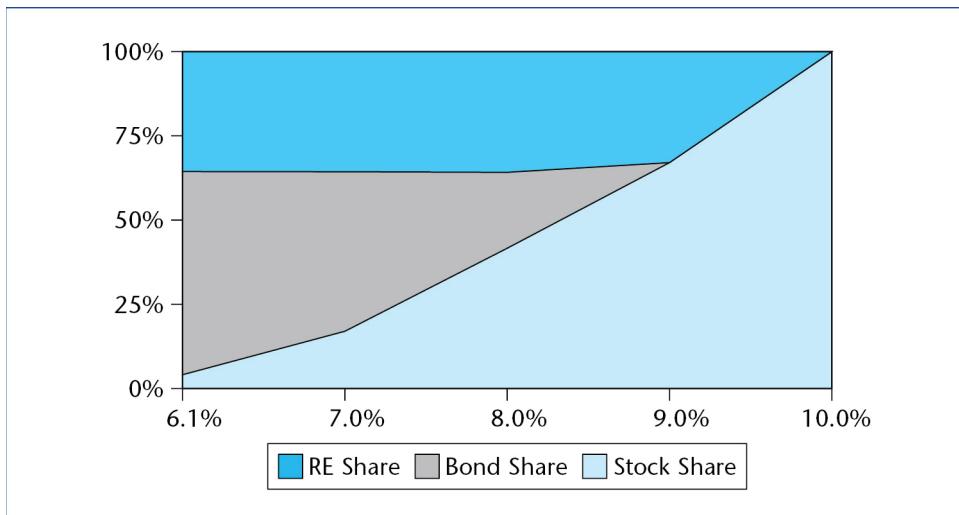


EXHIBIT 22-6 Asset Composition of the Efficient Frontier (based on Exhibit 22-3 expectations).

the optimal (efficient frontier) portfolio corresponding to the given target return where the vertical line intersects the horizontal axis.

As before, Exhibit 22-6 is based on the expectations in Exhibit 22-3. Notice that the optimal real estate share is rather large (over 30 percent) and stable for a broad range of conservative to moderately aggressive return targets (roughly from the minimum-variance target of just over 6 percent through at least 9 percent, which is not far below the maximum all-positive-weights portfolio return of 10 percent). The optimal share allocated to stocks rises steadily with the return target (i.e., with investor risk tolerance), first at the expense of bonds, and at more aggressive ranges at the expense of real estate. This makes intuitive sense, as real estate has a slightly higher expected return than bonds (though lower than stocks), and the correlation between real estate and stocks is assumed to be slightly lower than that between bonds and stocks. Thus, throughout the broad conservative and middle range of the risk preferences, real estate makes up about one-third of the optimal portfolio, while the balance of the optimal portfolio goes from mostly bonds and some stocks at the conservative end to mostly stocks and some bonds at the aggressive end, given the risk and return expectations of Exhibit 22-3.

The above optimal allocation results are an important empirical and investment policy relevant conclusion from the application of modern portfolio theory at the broad-brush level across the three major underlying asset classes, with obviously important implications for real estate investment. Although the specific results summarized in Exhibit 22-6 depend upon the specific risk and return expectations of Exhibit 22-3, those expectations are not implausible nor atypical, and similar optimal allocation results are obtained for other similar and similarly plausible expectations. In essence, if one divides the investment universe into only three underlying asset classes, then the role of each one of those asset classes will be generally fairly large, and the role of real estate, with its generally favorable risk/return expectations and low correlation to other asset classes, will tend to be greatest in the conservative to middle range of the risk/ return preferences.

This is a tremendously important conclusion for real estate investment at the big picture macro-level. It says that real estate should occupy a prominent role in the average or overall investment portfolio. Of course, it may seem that we are “stacking the deck” in favor of real estate by allowing for only three asset classes in our analysis and making real estate one of those classes. How could we not find a large role for real estate? But “big pictures” are just that: they present

the broadest outlines of reality. And the “biggest picture” in the investment world is, arguably, a three-asset-class world, with real estate one of those three. Real estate in fact underlies roughly one-third of the overall capital market value and represents close to one-third of the professionally investable universe of asset classes in the United States. One finds generally similar proportions in other countries. Stocks and bonds also each represent approximately similar proportions of the overall capital market value and the investable universe. Portfolio theory can provide an underlying explanation for this result!

TEXT BOX 22.2 IS MPT ECONOMICS?

We have noted that MPT is a famous theory that earned a Nobel Prize in economics for its developer, Harry Markowitz (1927–2023), and that it has been widely applied in the investment industry since the mid-twentieth century. But there are some interesting, unique characteristics of this theory in the world of economics. Recall that most of the fundamental theories and methods we have presented in this book are based on one or a combination of the three basic micro-economic principles and foundations: market equilibrium, opportunity cost, and wealth maximization. (Examples: The Quadrant Model is based on market equilibrium. The DCF valuation model is based on opportunity cost; and the NPV decision rule is based on wealth maximization.) Yet we have just presented the essence of MPT, and nowhere did we mention any of those three touchstones. Of course, the objective of maximizing return while minimizing volatility is certainly importantly related to wealth maximization, but it is not exactly the same thing. Indeed, Markowitz’s undergraduate degree at the University of Chicago was not in economics. When he was defending his PhD thesis on MPT, in the Chicago economics department, the famous economist Milton Friedman supposedly commented that it was a great theory, “but not economics.” (The story may not be apocryphal, as Markowitz personally told one of this text’s co-authors that many years later, when Markowitz was in a presentation with Milton Friedman in the audience, Markowitz turned to Friedman and said, “Now is it economics?”) Of course, we have already noted how the benefits of diversification violate the “no free lunch” stereotype of economic theories. Even before he won the Nobel Prize in economics, Markowitz won a prestigious prize in another field, operations research and management science, in part for his development of MPT. Whether it is “economics” or not, MPT is certainly a powerful and elegant tool for analyzing investment strategy. And we will see in the next chapter how a simplification of MPT can indeed be applied in an equilibrium model of the capital market to deepen our understanding of asset pricing and expected returns—the so-called Capital Asset Pricing Model (CAPM).

22.2 ALLOWING FOR A RISKLESS ASSET: THE TWO-FUND THEOREM AND THE SHARPE-MAXIMIZING PORTFOLIO

Portfolio theory as described in the preceding section provides an elegant and rigorous framework for thinking about strategic asset allocation at the level of the investor’s overall wealth portfolio. It provides insight and builds intuition relevant to all investors. It is widely used in practice, particularly by large institutional investors such as **pension funds**. However, an important consideration for many investors has not been included in the theory as presented so far. In particular, we have

considered only *risky* assets as possible components of the portfolio. By risky assets, we mean investments whose returns cannot be predicted with certainty in advance, investments whose periodic returns are volatile. Yet it is often useful to envision another type of asset, a *riskless* (or *risk-free*) asset.

In principle, a riskless asset's return is known in advance for certain, and it is viewed as having no volatility in the realization of its periodic returns. In reality, of course, there is no such thing as a perfectly riskless asset. However, the concept of the riskless asset is a useful construct because it provides an interesting extension of the theory described previously.

One reason for the introduction of a riskless asset in portfolio theory is that this construct can *approximate borrowing or short-term lending* by the investor. Borrowing, or the use of debt to finance investment, can be represented by a short position in the riskless asset.¹¹ When the investor borrows money (in effect *leveraging* her overall wealth, as described in Chapter 12), they normally intend to pay back the loan without defaulting. Thus, from the investor's perspective, the return on the debt asset appears certain, providing some logic for treating the debt as riskless. On the other hand, many investors hold a certain amount of "cash" in their wealth portfolios, that is, investments in very short-term debt. This is most commonly in the form of U.S. Treasury bills (T-bills), although private sector "commercial paper" and other instruments are also used. The periodic returns on such short-term lending by the investor have very little volatility and virtually no default risk, and so may be viewed as an approximately riskless investment.

22.2.1 Two-Fund Theorem

While the riskless asset construct is useful as an approximation for borrowing and short-term lending, it has another, more technical use as well in portfolio theory. The riskless asset construct can be used as a sort of mathematical technique to greatly simplify the portfolio allocation problem. This simplification is known as the **two-fund theorem**, and it, too, is widely used in practice.

To see what the two-fund theorem means and how it works, you need to understand a basic mathematical fact. Recall from our discussion of Exhibit 22-4B that whenever two risky assets are less than perfectly positively correlated, the risk/return possibilities of combinations of those two assets lie along a curve that is bent upward and leftward in the risk/return diagram. But now suppose one of the two assets is riskless, that is, contains no volatility. In this case, the risk/return possibilities of combinations of the two assets will lie exactly on a straight line connecting the risk/return possibilities of each of the two assets alone.

This means that, in a risk/return diagram, the risk/return possibilities from a combination of a riskless asset and *any* portfolio of risky assets will lie along a straight line connecting the riskless asset's risk/return with the risky portfolio's risk/return. This is depicted in Exhibit 22-7A for combinations of a riskless asset with an interest rate of r_f and a risky portfolio with risk and expected return represented by point Q . All the risk/return possibilities lying on the straight line connecting r_f with Q are feasible by holding the portfolio Q combined with either long or short positions in the riskless asset. If the investor puts some of their wealth in the portfolio Q , and some in the riskless asset, then their risk and return will lie on the line segment between points r_f and Q , to the left of Q .¹²

¹¹ Recall that a short position can be represented by a negative investment allocation. In the case of short positions, the investor receives cash when the position is taken and gives up cash when the position is closed or retired. This is in direct contrast to long positions, which are characterized by the fact that cash is given up at the time the investment position is taken, and cash is obtained when the asset is sold to close out the position. Thus, you can see how borrowing can be represented as a short position in a debt asset, just as lending is a long position in such an asset.

¹² This would be short-term lending, that is, taking a long position in the riskless asset.

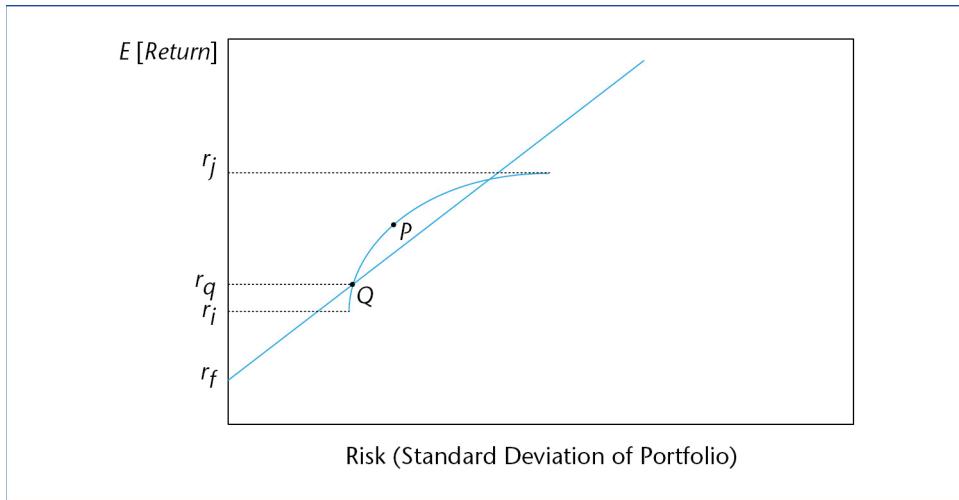


EXHIBIT 22-7A Risk and Return Possibilities of a Combination of a Riskless Asset with Return r_f and the Suboptimal Risky Portfolio with Expected Return r_q .

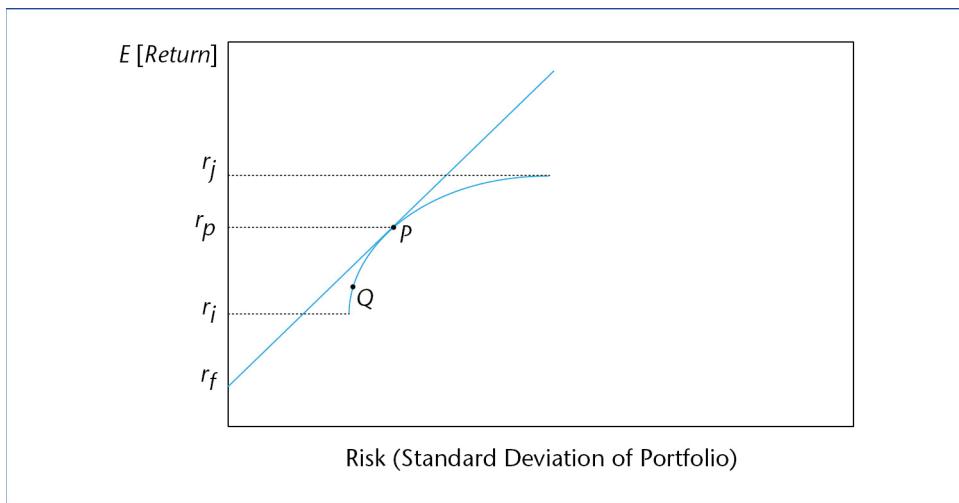


EXHIBIT 22-7B Risk and Return Possibilities of a Combination of a Riskless Asset with Return r_f and the Optimal Risky Portfolio with Expected Return r_p .

(The greater the share of wealth in the riskless asset, the closer to r_f their expected return will be, and the farther to the left of Q on the straight line their risk and return will lie.) If the investor borrows money, then their risk and return will lie on the straight line to the right of Q .¹³

¹³ This would be taking a short position (placing a negative proportion of the investor's wealth) in the riskless asset. As the proportion of the investor's wealth in both the riskless asset and the risky portfolio must sum to unity, this implies that the proportion of the investor's wealth placed in the risky portfolio Q would exceed unity. In other words, borrowing enables the investor to place more than their total wealth in the risky portfolio, their own net wealth plus the amount they have borrowed against it. As the expected return to Q exceeds the risk-free rate of r_f at which they are borrowing, the resulting leverage is "positive" (to use the terminology introduced in Chapter 12), causing the investor's expected return to exceed r_q , lying therefore to the right of Q on the straight line connecting r_f and Q .

In Exhibit 22-7A, the curved line represents the efficient frontier of risky assets described in Section 22.1.3. Thus, the risky portfolio Q is an efficient portfolio, in the absence of the possibility of investing in a riskless asset. Nevertheless, if it is possible to take positions in a riskless asset, no investor would want to combine such positions with Q . All of the risk/return possibilities on the straight line connecting r_f with Q are dominated by possibilities lying on the straight line connecting r_f with P , a different portfolio also lying on the efficient frontier of risky assets. This is depicted in Exhibit 22-7B. In fact, you can see by geometric reasoning that the possibilities on the straight line connecting r_f and P dominate any and all other feasible risk/return possibilities. P lies at the point on the risky asset efficient frontier (the curved line) that is just tangent to a straight line passing through r_f ¹⁴

This brings us to the two-fund theorem. According to this theorem, all investors (no matter what their risk preferences) will prefer combinations of the riskless asset and *a single* particular risky asset portfolio (the point P in Exhibit 22-7B). With the possibility of riskless borrowing or lending, the optimal allocation among risky assets, P , is determined not by the investor's risk preferences, but only by their expectations about asset future returns. The investor's risk preferences can be met by adjusting the position in the riskless asset only, by either borrowing more (for more aggressive preferences) or lending more (for more conservative tastes), so as to meet the appropriate expected return target.

22.2.2 SHARPE-MAXIMIZING PORTFOLIO

How can we determine which combination of risky assets corresponds to the unique optimal allocation represented by point P in the two-fund theorem? The answer to this question is obvious in the geometry we just used to discover the theorem. The point P is the portfolio that maximizes the slope of a straight line connecting that portfolio's risk and return with the risk and return of the riskless asset. Of course, the riskless asset has zero risk, and its return is r_f . The risk (volatility) of the risky-asset portfolio is S_p , and its return is r_p . The slope of the line connecting these two points is thus: $(r_p - r_f) / S_p$. This is the risk premium of the risky asset portfolio divided by its volatility. This ratio is known as the **Sharpe ratio**, named after William F. Sharpe, who shared the Nobel Prize for portfolio theory. In the presence of a riskless asset, the optimal combination of risky assets is the one with the highest Sharpe ratio.¹⁵

While the two-fund theorem has an elegance that appeals to the aesthetic sense of academicians, it also has a simplicity and intuitive appeal that practitioners like. It is thus widely used in practice. As the Sharpe ratio is just the risk premium (what we have been labeling RP) divided by the risk (defined as volatility), it is a natural measure of **risk-adjusted return** and is widely (and somewhat

¹⁴ Any line northwest of P would not be feasible, as it would pass entirely through the infeasible region to the northwest of the risky asset efficient frontier. In other words, it would be impossible to find a risky portfolio to combine with the riskless asset to produce risk/return possibilities that lie on such a line. Any straight line through r_f that would run southeast of P would be dominated by the line through P .

¹⁵ You can use the solver on a spreadsheet to find the Sharpe-maximizing allocation the same way as is described in the demonstration video in the book's Online Supplemental Materials (see www.routledge.com/cw/geltner-miller) to find the variance-minimizing allocation. Simply compute the Sharpe ratio of the portfolio in a cell of the spreadsheet using formulas A.1 and A.6 from Appendix 22 (as in the demo Excel file in the Supplemental Materials). Then, instead of telling the solver to minimize the portfolio variance, tell it to maximize the portfolio Sharpe ratio. You will no longer need the target return constraint in the solver. (The resulting optimal allocation will just be for the optimal risky asset portfolio, P . You will have to lever this portfolio up or down using the riskless asset to meet a given target return.) Obviously, somewhere in the spreadsheet, you have to input the risk-free rate, r_f . The Sharpe-maximizing allocation will be different for different values of r_f .

informally) used as such by practitioners. The numerator of the Sharpe ratio measures the investor's compensation for risk, the excess return over what an investment in Treasury bills would provide. The denominator of the Sharpe ratio is the amount of risk in the investment as measured by the standard deviation of its periodic returns. If the Sharpe ratio is applied at the level of the investor's entire wealth portfolio, volatility is arguably the relevant measure of risk. Furthermore, the Sharpe ratio is unaffected by the investor's position in the riskless asset, either by borrowing or lending risklessly. Thus, it makes intuitive sense to want to find the risky asset allocation that maximizes the Sharpe ratio.

$$\text{Sharpe ratio of } P = \frac{r_p - r_f}{S_p} = \frac{RP_p}{S_p}$$

In considering the role of real estate in the optimal mixed-asset portfolio, the Sharpe-maximizing perspective sometimes reduces the real estate share for conservative investors, as compared to the perspective from classical MPT without a riskless asset. In general, the Sharpe-maximizing risky asset allocation tends to place lower weights in low-return asset classes, as compared to the variance-minimizing allocation with a conservative target return.

As an example, consider again the risk and return expectations in Exhibit 22-3. To meet a (rather conservative) 7 percent return target, recall that the variance-minimizing shares among the three risky asset classes were: 16 percent stocks, 48 percent bonds, and 36 percent real estate. Now suppose we can approximate the risk-free interest rate as being 3 percent. The Sharpe-maximizing risky asset allocation (P) would then consist of 27 percent stocks, 37 percent bonds, and 36 percent real estate. The relative proportion of stocks is considerably greater, and long-term bonds smaller, than in the 7 percent target variance-minimizing portfolio. However, this comparison is a bit misleading because this all-risky-asset Sharpe-maximizing portfolio is slightly too aggressive to meet the risk level implied by the 7 percent target. This is seen by the fact that, without any investment in T-bills, the **Sharpe-maximizing portfolio** would have an expected return of 7.4 percent with 7.49 percent volatility (compared to 6.89 percent volatility in the 7 percent target all-risky-asset variance-minimizing portfolio).

EXHIBIT 22-8

Comparison of Optimal 7%-Return-Target Portfolio Allocations, Variance Minimization Versus Sharpe Ratio Maximization

	Return and Risk Expectations*			Portfolio Allocations	
	Return	Volatility	Sharpe Ratio	Var.-Min.	Sharpe-Max.
Cash (T-bills)	3.00%	NA**	NA	NA	10%
Bonds	6.00%	8.00%	0.38	48%	33%
Real Estate	7.00%	10.00%	0.40	36%	32%
Stocks	10.00%	15.00%	0.47	16%	25%
Var.-Min. Portfolio	7.00%	6.89%	0.58	100%	NA
Sharpe-Max. Portfolio	7.00%	NA**	0.59	NA	100%

* Also includes correlations:

Stock/Bond +30%, Stock/Real Estate +25%, Bond/Real Estate +15%.

** From the Sharpe-maximization perspective, T-bills are viewed as having zero volatility, but as this is not exactly true in reality, it would be misleading to calculate and show a Sharpe-maximizing portfolio volatility juxtaposed with that of the variance-minimized portfolio.

To meet the 7 percent return target using the Sharpe-maximizing portfolio, we would have to allocate 90 percent of the investor's wealth to the Sharpe-maximizing risky asset portfolio and 10 percent to T-bills. The overall allocation would then be 25 percent stocks, 33 percent long-term bonds, 32 percent real estate, and 10 percent T-bills (cash). The combined long- and short-term debt allocation (bonds and bills) would be 43 percent, slightly less than the 48 percent allocation to long-term debt alone in the variance-minimizing 7 percent target portfolio. Thus, the optimal 7 percent target portfolio from a Sharpe-maximizing perspective allocates considerably more weight to stocks (25 percent instead of 16 percent) and slightly less weight to both bonds (43 percent instead of 48 percent) and real estate (32 percent instead of 36 percent). In a sense, part of the reason the bond and real estate allocations are so high in the variance-minimizing portfolio is simply because the investor has specified a conservative return target, and real estate has a relatively low expected return. These differences between the Sharpe-maximizing and variance-minimizing perspectives are summarized in the table in Exhibit 22-8.¹⁶ Note that, even though the optimal real estate allocation is reduced in the Sharpe-maximizing portfolio, it still retains a large share, given the indicated expectations.

22.2.3 SUMMARY OF THE IMPLICATIONS OF THE RISKLESS ASSET CONSTRUCT

In this section, we have seen how the extension of the original Markowitz portfolio theory to include consideration of a riskless asset can be useful in several ways:

- It allows an alternative, intuitively appealing definition of the optimal risky asset portfolio, the one with the maximum Sharpe ratio.
- It can help to avoid "silly" portfolio recommendations that put too little weight in high-return assets just because the investor has a conservative target return (or vice versa).
- It provides a useful framework for accommodating the possible use of leverage or cash in the portfolio.

On the other hand, the riskless asset construct is an extension of the original model that adds an additional, not-quite-realistic assumption. The use of one form or the other of the portfolio model depends in part on taste, and in part on circumstances. In practice, both versions are widely used in the real world.

22.3 OTHER MODELS OF PORTFOLIO ALLOCATION: RISK PARITY AND PASSIVE ALLOCATION

The preceding two sections introduced you to classical portfolio theory. Either with or without the riskless asset construct, this is the most widely taught conceptual tool for considering macro-level investment allocation decisions in the framework of risk and return. It is an elegant and powerful theory for developing strategic insight about investment allocation. It is most widely taught because of its rigor and ability to build your intuition and common sense about how to think about portfolio allocation. But while MPT is the classical and standard approach, we noted at the outset of this chapter that two other approaches are widely used among large investment institutions. These two approaches, risk parity and passive allocation, may often play a larger role in the final portfolio allocation decision, even though they are less grounded in formal theory and do not involve explicit optimization.

¹⁶ In Exhibit 22-8, T-bills are used as a proxy for the risk-free asset. As an example of the computation of the Sharpe ratio, the real estate ratio is computed as $(7\% - 3\%)/10\% = 0.40$. Note that the Sharpe ratio of a portfolio is unaffected by its allocation to the risk-free asset.

The first such approach is what is called **Risk Parity**, which involves the identification of the “Risk Parity Portfolio” (RPP). The RPP is not exactly an “optimum”, and it is not based directly on any rigorous economic model. But it is appealing, including for some valid economic considerations. The RPP is the portfolio allocation that causes each asset class to have the same contribution to the volatility of the portfolio. With the RPP, the investor is exposed to an “equal amount of risk” from each of the constituent asset classes (where “risk” is defined as the portfolio volatility). This is a sort of intuitive way of thinking about a “diversified” portfolio. The concept of the RPP is attributed to the investment firm Bridgewater, in the 1990s. Provided the RPP is not very far off of (that is, below) the mean-variance efficient frontier, there can be a valid economic argument in favor of the RPP.¹⁷ It goes something like the following.

While investors are averse to “risk”, they are even more averse to “uncertainty.” The difference between “risk” and “uncertainty” is that “risk” is quantifiable. With risk, you can know (or reasonably assume) the relevant probability of future outcomes. With “uncertainty”, you “don’t know what you don’t know.” You can’t quantify the probabilities. Sometimes people refer to “risk” as “known unknowns” and “uncertainty” as “unknown unknowns.” With this in mind, consider that the outputs from a model are no better than the quality of the inputs into the model. The Markowitz model is great in theory, but it requires input data in the form of the expected future returns to the asset classes (as well as the volatilities and correlations). But it is very difficult to predict future returns to asset classes. There is likely a great deal of uncertainty about the means of the returns of the asset classes, not in past history, but in what actually matters for decision making, which is in the forward-looking future expectations. The RPP is robust to this type of uncertainty, because in it the share of each asset class in the portfolio does not depend on the means of the returns, the expectations, the RPP depends only on the asset second moments (the volatilities and correlations), which are what determine the volatility of the portfolio. For a given set of second moment expectations, there is only one Risk Parity Portfolio. It is not a whole curve of points like the MPT Efficient Frontier. It’s only one point in the risk/return diagram. And this simplicity is also appealing from a practical perspective.

The other less rigorous but very popular approach to portfolio allocation is what is often called **“passive allocation”**, or **“managing to the benchmark.”** The idea here is that “the market knows best.” You observe the value shares of the various asset classes that comprise the market of all of the relevant assets, and you match that allocation in your own portfolio. At the broadest level, the relevant market would be the national investable wealth portfolio, consisting of all real estate, stocks, bonds, and other relevant asset classes. For example, if we take the asset class shares described in Exhibit 1-4 back in Chapter 1, we see that as of around 2020, all forms of real estate made up at least 36 percent of the total investable capital market “pie” in the United States. This would suggest a target allocation to real estate in your portfolio of a similar share, about 36 percent. If you want to consider only commercial real estate equity, Exhibit 1-4 would suggest a real estate share in the mixed-asset portfolio of more like 10 percent. In Chapter 23, we will describe the theory of the classical Capital Asset Pricing Model (CAPM), which can provide something of an intellectual basis for the passive allocation approach. From a practical perspective, passive allocation is also relatively simple and intuitive. It tends to require less research and resources than approaches such as MPT or RPP that require quantitative estimates of return expectations.

Often, investment managers face a “benchmark” defined by a quantifiable, regularly reported investment performance index based on a basket of investment assets or products. The manager’s performance will be judged relative to the performance of that benchmark. In that case, a “passive” investment strategy would be to not try to beat the benchmark, but rather reduce the risk of deviating

¹⁷ By “not very far off” the efficient frontier, we mean that the expected return and volatility of the RPP are very similar to that of some point on the efficient frontier. The RPP cannot possibly be above the efficient frontier, since the frontier is defined by the maximum expected returns possible for each level of portfolio volatility.

from the benchmark by matching your portfolio allocation to that reflected in the benchmark index. This is another way to define the “passive allocation” approach in practice. It involves not “going out on a limb” if you are an investment manager. But it tends to result in “herd behavior” in the investment industry.

22.4 CHAPTER SUMMARY

This chapter introduced you to the most famous and fundamental tool in macro-level real estate investment analysis, modern portfolio theory. This theory, including its extension using the riskless asset construct, provides the core discipline for rigorous strategic investment decision making. This is the level of decision making relevant to deciding broad, long-run capital allocation policies across the major asset classes. As real estate is one of the major asset classes, the perspective, intuition, and methodology we presented in this chapter are of vital importance to anyone seriously interested in real estate investment. Finally, we also briefly described two other less rigorous but very popular portfolio allocation approaches that are widely used at the mixed-asset portfolio level where decisions are made about how much to invest in real estate. It is important to note that by any and all of the three portfolio allocation approaches described in this chapter, the real estate asset class is generally found to merit a significant allocation in the mixed-asset portfolio.

KEY TERMS

- Modern (or Markowitz, or Mean-variance) Portfolio Theory (MPT)
- Volatility (standard deviation of return)
- Covariances (between assets)
- Weighted covariances
- Correlation coefficient (between returns)
- Solver (in spreadsheet) target return (for portfolio)
- Sensitivity analysis
- Mean-variance portfolio theory (Markowitz)
- Risk preferences (of investors)
- Dominant portfolio (domination)
- Diversification benefit
- Cost-of-living adjustments (COLAs)
- Mixed-asset portfolio
- Efficient portfolio
- Efficient frontier
- Pension funds
- Two-fund theorem
- Sharpe ratio
- Risk-adjusted return
- Sharpe-maximizing portfolio
- Informational efficiency
- Illiquidity
- Risk Parity Portfoio
- Passive allocation
- Managing to the benchmark
- Risk vs Uncertainty

STUDY QUESTIONS

Conceptual Questions

22.1 Why should all investors, regardless of risk preferences, avoid holding *dominated* portfolios?

22.2 What does modern portfolio theory recommend besides not putting all one's eggs in one basket?

22.3 Why do investors invest in real estate? Answer this from a portfolio theory perspective, considering the trade-off between portfolio risk and return. (Hint: What is the primary role of typical income property real estate within the portfolio? Is it bought to achieve a high average return?)

22.4 Building off of your answer to Question 22.3, what is another role for real estate in helping to manage the portfolio of a defined benefit pension fund portfolio?

22.5 How are investor risk preferences reflected in typical portfolio theory application? (Hint: What is the role of the target return specified by the investor?)

22.6 Modern portfolio theory tells us that investors hold portfolios that lie on the efficient frontier and that efficient frontier portfolios change with investor's target required portfolio returns. Examine the role of real estate in optimal portfolios with different risk-return profiles by answering the following question:
Explain how the optimal portfolio allocation or mix changes as we move from the more conservative to aggressive investor return expectations in a world with three major investment asset classes (stocks, bonds, and real estate) as shown in Exhibit 22-6.

22.7 How is the two-fund theorem useful in practical portfolio theory applications? What is the relationship between the two-fund theorem and the Sharpe-maximizing portfolio?

22.8 Available empirical evidence suggests that commercial real estate is an asset class that should be well represented in the mixed asset portfolios of institutional investors. However, despite the perceived importance of real estate in mixed asset portfolios, pension funds typically allocate smaller shares to real estate than analysis based on Modern Portfolio Theory (MPT) predicts they should. Discuss why this might be so.

Quantitative Problems

22.9 You work for BSC Advisors, a major real estate investment manager. You are trying to convince a pension fund manager to allocate 10% of her portfolio to investment-grade income-property. Currently the fund is allocated 60% to stocks and 40% to bonds. Suppose you agree with the asset class risk and return expectations shown in Exhibit 22-3. Perform the following calculations, and then explain how you would use the results to persuade the pension fund portfolio manager to allocate funds to real estate. Be prepared to discuss the limitations of your analysis. (Help: A portfolio comprised of three asset classes ("1", "2", and "3") has expected return, $E[R_p]$, and standard deviation of returns (volatility), S_p , given by

$$E[R_p] = w_1 E[R_1] + w_2 E[R_2] + w_3 E[R_3] \text{ with } w_1 + w_2 + w_3 = 1 \text{ and}$$

$$S_p = \sqrt{w_1^2 S_1^2 + w_2^2 S_2^2 + w_3^2 S_3^2 + 2w_1 w_2 S_1 S_2 C_{12} + 2w_1 w_3 S_1 S_3 C_{13} + 2w_2 w_3 S_2 S_3 C_{23}}$$

where the w_i are the portfolio weights for each asset class and the C_{ij} are the correlation coefficients between pairs of asset class returns.)

- a. Determine the expected return and standard deviation of returns on the current portfolio (60% stocks/40% bonds).
- b. Determine the expected return and standard deviation of returns on a portfolio comprised of 55% stocks, 35% bonds and 10% real estate.

22.10 Consider two portfolios. Portfolio A has an expected return of 10% and volatility of 8%. Portfolio B has an expected return of 9% and volatility of 7%. The interest rate on a risk-free investment is 5%. Which of the two risky portfolios is not on the efficient frontier? (Hint: Use the two-fund theorem.)

22.11 Consider again portfolios A and B from Problem 22.10. What would be the expected return to portfolio C that consists of a 50/50 weighting of A and B?

22.12 Consider again portfolio C from Problem 22.11. Assuming (very reasonably) that portfolios A and B are not perfectly positively correlated, will portfolio C's volatility be greater or less than halfway between A's and B's volatility (that is, 7.5%)?

23 Equilibrium Asset Valuation and Real Estate's Price of Risk in the Capital Market

CHAPTER OUTLINE

- 23.1 Introduction and Some Threshold Points
 - 23.1.1 Practical Uses for Asset Pricing Theory
 - 23.1.2 A Threshold Point: What Underlies Asset Risk?
- 23.2 Review of Classical Asset Pricing Theory
 - 23.2.1 From Portfolio Theory to the CAPM
 - 23.2.2 The Main Point in the Basic CAPM
 - 23.2.3 Isn't the CAPM Wrong?
 - 23.2.4 Strengths and Weaknesses in the Basic CAPM
- 23.3 Applying the CAPM to the Private Real Estate Asset Class as a Whole
- 23.4 Attempting to Quantify Risk and Return within the Private Real Estate Asset Class
- 23.5 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- What is meant by equilibrium asset pricing models and how these tools are used in practice to aid macro-level investment decision making.
- The classical Capital Asset Pricing Model (CAPM) and its major theoretical and practical strengths and weaknesses.
- How the CAPM can be (or why it should be) applied to real estate and where it falls short in this regard, including a distinction between application at the overall multi-asset-class portfolio level and application at the more specific level within the private real estate asset class.

The previous chapter introduced you to portfolio theory, the fundamental tool of **strategic investment decision making** at the macro-level. This chapter will build on portfolio theory to present a topic that is relevant and useful in principle for all three of the decision-making arenas noted in the Introduction to Part VII: tactical as well as strategic investment decision making, plus investment policy implementation. Our subject in this chapter will be equilibrium asset pricing models. Such models are typically built on the foundation of portfolio theory, but go beyond portfolio theory to provide a simplified representation of how the capital market perceives and prices risk (or other attributes of concern to investors) in the assets that are traded in the market. Intuitively, we know that assets that are riskier (in some sense that the capital market cares about) will trade in equilibrium at lower prices, so that they will provide higher expected returns.¹ Thus, asset pricing models at the

¹ Recall the inverse relationship between asset value and expected return. See, for example, Chapter 4.

macro-level are, in effect, models of the ex-ante risk premium required by investors.² This means they can, in principle, be used to help forecast future long-run average returns.

This chapter will be presented in four parts. First, we provide some introductory and threshold points in Section 23.1. Then in Section 23.2, we introduce the classical equilibrium asset pricing model in mainstream securities investments, the **capital asset pricing model (CAPM)**. Section 23.3 applies this model at the overall level of the mixed asset class portfolio, stocks, bonds, and real estate. Finally, Section 23.4 explores how, or whether, such modeling can be applied at a more granular level within the real estate asset class.

23.1 INTRODUCTION AND SOME THRESHOLD POINTS

Before delving into the nitty-gritty of equilibrium asset price modeling itself, let us put this exercise in some perspective. This section will review the practical uses of the models, and then raise an important threshold point about the nature of the risk that is being modeled.

23.1.1 PRACTICAL USES FOR ASSET PRICING THEORY

As noted, equilibrium asset pricing models have three major practical uses. First, asset valuation and tactical investment decision making are naturally linked because **tactical decision-making** aims at profiting from short- to medium-term opportunities. This often involves identifying assets (or, more likely, asset classes, or *types* of assets or investment products) that are currently *mispiced* in some sense, such that there might be short- to medium-term opportunities in buying or selling. In order to know if something is mispriced, you need to have some idea about what its value should be when asset markets are in equilibrium. Asset pricing models give insight about what those equilibrium asset values and expected returns should be and thus, where they are headed.

Equilibrium pricing models can also be useful for strategic investment decisions, that is, for long run or permanent objectives for a given investor, in the following way. Suppose the model can identify a risk factor that is persistently and reliably *priced* in the capital market, but that a given investor is not concerned about. For example, say the capital market prices smaller properties in tertiary markets as though they are relatively more *risky*. In other words, such investments tend to provide an extra total return premium just because they are small and in tertiary markets. This would presumably be because the *average* investor (determining the equilibrium in the asset market) is persistently concerned about the particular risk in such assets (or their illiquidity, or some other negative attribute like the difficulty of getting there). If you as an investor, don't share this concern that the average investor has, then the model has just given you a strategic reason to tilt your portfolio allocation a little bit extra towards smaller properties in tertiary markets. You would stand to earn an extra return over the long run (not *definitely* but *probably*, if the model is correct and the equilibrium is persistent) only by taking on an extra risk factor that you are not concerned about (in contrast to the average investor).

Finally, equilibrium asset price models provide a rigorous way to quantitatively adjust realized investment performance (ex-post returns) for the amount of risk the investment was exposed to, based on the pricing of such risk in the capital market. In principle, this can enable a more fair and complete or correct evaluation of the performance of investment managers hired to implement investors' policies (including regarding risk exposure). If an investment manager did very well, but she did so by taking on a lot of risk, then perhaps a big part (or even all) of the high return she

² To put macro-level valuation in a broader context, note that the most widely used asset valuation model at the micro-level is the DCF model described extensively in Parts I–III. There, we noted the importance of the expected return risk premium in deriving the discount rate used in the denominators of the DCF valuation equation, and we noted that this risk premium reflects the opportunity cost of capital (OCC) as this is determined ultimately in the capital market. However, we stopped short of considering in depth *how* the capital market determines the expected return risk premium in the OCC. This is the question we turn to in the present chapter.

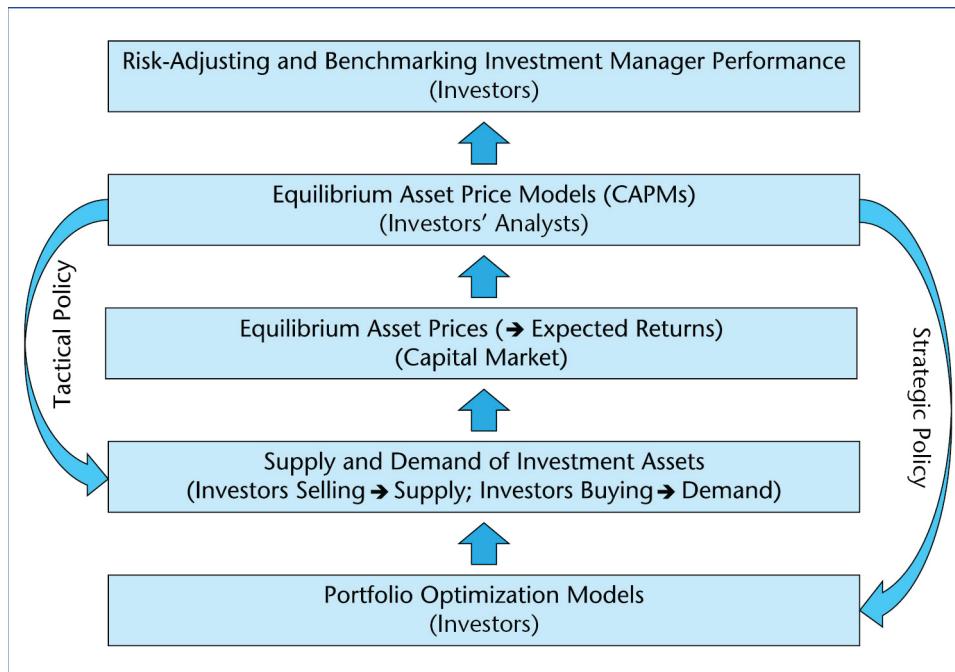


EXHIBIT 23-1 The Relationship between Equilibrium Asset Price Models and Investment Policy.

achieved was simply a compensation for risk exposure. Perhaps she wasn't "beating the market." Thus, asset pricing models can help to *benchmark* investment performance.³

The context and roles of equilibrium asset price models are illustrated in Exhibit 23-1. Underlying investors' decisions are portfolio optimization models, a classical example of which (the MPT—Markowitz Mean-Variance Model) was described in the previous chapter. Investors' decisions (which assets to buy and sell) lead to the pattern of equilibrium asset prices in the capital market. This equilibrium pricing is the subject of the asset price models (such as the classical "CAPM," which we will describe shortly) that seek to specify what are the equilibrium (or in some sense "fair market") expected returns to the various assets or asset classes. This information can then be used by investors: (1) to help inform their long-run strategic portfolio optimization policies due to differences between their preferences and the average market preferences (reflected in the equilibrium); (2) to help identify short to medium-term tactical opportunities due to mispricing in the market; and (3) to help evaluate or benchmark the performance of investment managers hired to implement the investors' strategic or tactical policies.

23.1.2 A THRESHOLD POINT: WHAT UNDERLIES ASSET RISK?

Before we examine how to model the way asset risk is priced, let us step back briefly to consider the nature of that risk. The risk that matters to an investor in an asset is based on the risk in the *total return* that the asset provides. That is, what matters to the investor is both the income provided by the asset and the change over time in the value of the asset. Fundamentally, therefore, there are two sources of investment risk: **cash flow risk**, and **asset-valuation risk**. In the case of real estate assets, cash flow risk derives from changes over time in the space market that cause changes in the rent the property

³ This last topic will be addressed in Chapter 25, after we have covered more of the elements of macro-level real estate investment analysis.

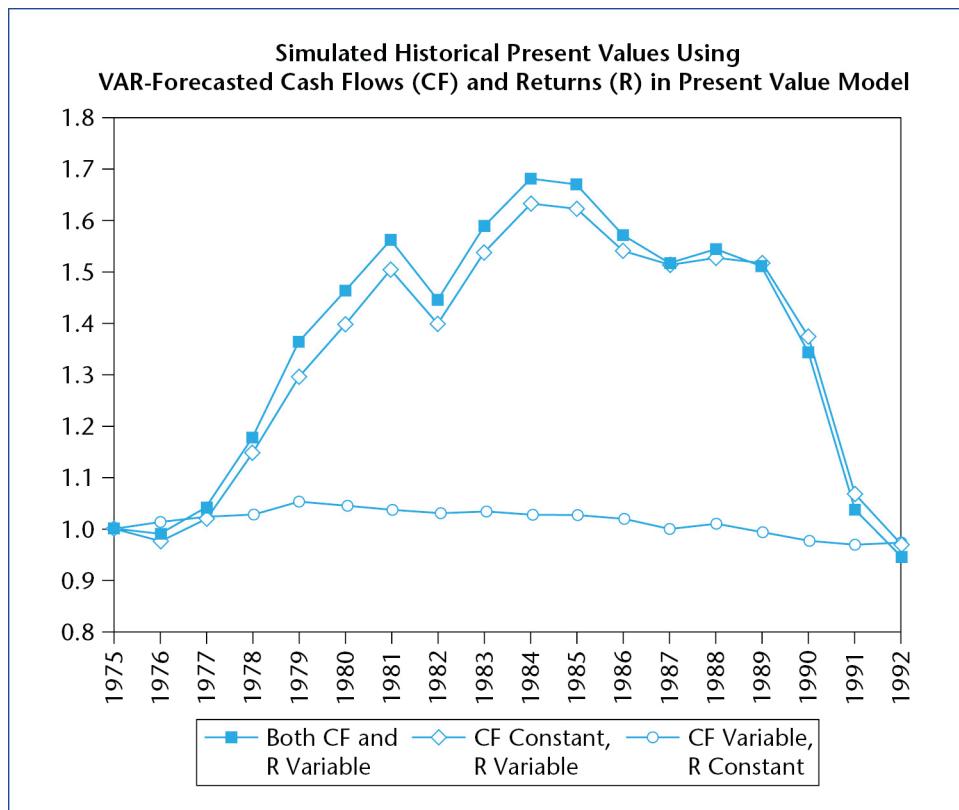


EXHIBIT 23-2 The Effect of Varying Cash Flow Expectations and Varying Return Expectations in a Present Value Model of Commercial Property.

Source: Geltner and Mei (1995).

can charge, as well as changes in operating expenses and capital improvement expenditures that also determine net cash flow. Asset-valuation risk reflects changes over time in the asset market that cause changes in the opportunity cost of capital applied by the asset market (the *discount rate* in the DCF procedure of Chapter 4), that is, changes in the equilibrium expected return required

While the space market and the asset's net cash flows are the fundamental source of property value, there is evidence that investment return risk in real estate assets, like that in the stock market, derives at least as much from changes in the asset market. Time variation in the discount rate may cause at least as much volatility in real estate investment returns as time variation in the property cash flow. This phenomenon in the stock market was first noted by the economist Robert Shiller in a famous 1981 article:⁴ “Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?” Shiller’s answer was “yes,” and although subsequent research has debated whether the word “*Justified*” in Shiller’s title is, itself, *justified* (see Section 23.4), there seems little doubt that a substantial part of the variation in asset prices is due to *revaluation* in the asset market, rather than to investors’ revisions in their expected future cash flows.

A picture of this point is presented in Exhibit 23-2 (taken from research by Jianping Mei and David Geltner⁵). The exhibit shows three lines. The line with solid square markers is the history of real estate

⁴ Shiller, R. 1981. “Do stock prices move too much to be justified by subsequent changes in dividends?” *American Economic Review* 71(3): 421–436.

⁵ Geltner, D. and J. Mei. 1995. “The present value model with time-varying discount rates: Implications for commercial property valuation and investment decisions,” *Journal of Real Estate Finance & Economics* 11(2): 119–136.

present values based on a DCF model in which changes in both the future cash flows and the future discount rate are forecasted and used in the valuation.⁶ This line closely tracks the actual market values of commercial property as represented by the NCREIF Index (unsmoothed). The similar line with diamond markers is the same DCF valuation model, only assuming a constant future cash flow level from each period going forward but including the forecasted time-varying discount rate in the valuation. The nearly flat line at the bottom, which does not much resemble what actually happened to commercial property values, applies the DCF model with a constant discount rate, but using the forecasted cash flows. While there was substantial variation in real estate cash flows, much of this variation was mean-reverting and predictable. Hence, the cash flow forecast does not change much from one period to the next, resulting in little change in asset present value unless the discount rate changes.

Why is an understanding of this fundamental nature of real estate investment return risk important as you proceed to consider how that risk is priced? As you will see, the essence of the CAPM is that priced risk is not the isolated volatility of individual assets, but rather is related to the co-movements of asset returns. When you realize that much asset value volatility derives from asset market valuation changes, you can appreciate how important asset markets, as well as space markets, can influence those all-important co-movements that govern the risk that matters. For example, it is possible for real estate returns across different space markets to move together much more than the underlying space markets themselves may tend to move together. With this in mind, let us proceed to a presentation of the CAPM.

23.2 REVIEW OF CLASSICAL ASSET PRICING THEORY

In this section we will introduce and briefly review the classical equilibrium asset pricing model, **CAPM (capital asset pricing model)**. It is such a famous and fundamental model that it is known traditionally simply as *the CAPM*.

23.2.1 FROM PORTFOLIO THEORY TO THE CAPM

The so-called Sharpe-Lintner CAPM is one of the most famous theories in all of financial economics (having earned William Sharpe and John Lintner a Nobel Prize, in 1990, together with Harry Markowitz). The CAPM grew out of, and is built on, the Markowitz mean-variance portfolio theory (MPT) that we described in the previous chapter. So, let's pick up where we left off in that chapter and see, in a simplified intuitive way, how we can extend MPT to develop a theory about how all assets in the market are priced. We can do this in four easy steps, as follows.⁷

Step 1: The Two-Fund Theorem. We begin with the basic assumptions of MPT, that investors want to maximize return and minimize risk in their wealth portfolios (i.e., they don't want to hold dominated portfolios), and we add the assumption described in Section 22.2 of Chapter 22 that a *riskless asset* exists. Recall that this leads to the two-fund theorem, which demonstrates that *all* investors (no matter what their risk preferences) should want to hold *the same* portfolio (i.e., the same relative weights) of risky assets, as long as those investors have the same risk and return expectations.

Step 2: Common Expectations. Remember that asset markets (especially securities markets such as the stock market) are known to be pretty efficient, that is, asset prices are pretty good

⁶ In effect, instead of the constant "blended" discount rate presented in Chapter 4, the DCF valuation here forecasts the future OCC of real estate in each future period of time, and applies those future forecasted rates to discount the forecasted future cash flows in each period. For example, the PV of the cash flow two years hence would be, in essence: $E[CF2]/((1 + E[r_1])(1 + E[r_2]))$, where r_1 is the OCC in year 1 and r_2 is the OCC in year 2, and $E[-]$ represents forecasted values as of the present. (The actual model used in the study is a log-transformation.)

⁷ The CAPM can be derived in several different ways, depending on what simplifying assumptions one wishes to make. The approach presented here corresponds closely in spirit (although not in analytical rigor) to the original development of the model.

at reflecting all publicly available information relevant to their values. Thus, any one investor will not usually have better information than the market as a whole concerning the values of assets or their future expected returns. Thus, everyone will converge to having the same expectations regarding risks, returns, and values. Combined with the two-fund theorem, this will lead everyone to want to hold the same portfolio (relative weights) of risky assets, which will therefore be empirically observable as the **market portfolio**, the aggregate of all risky capital assets. The relative weight placed on each asset (or asset class) in this market portfolio corresponds to the current market value of each asset as a fraction of the aggregate market value. This market portfolio thus reflects the overall wealth portfolios of all investors.⁸

Step 3: Only Covariance with the Market Matters. Now recall that a basic tenet of portfolio theory is that all investors care only about the risk (variance or volatility) in their *overall* wealth. The variance of any individual component of that wealth does not matter per se. Since, by our reasoning in steps 1 and 2, all investors hold the same portfolio, the market portfolio, the only risk that matters to investors in any given asset is how that asset affects the risk in the market portfolio. But it is a fact of basic statistics that the marginal contribution of each asset to the variance of a portfolio is directly proportional to the covariance of the asset's periodic returns with the portfolio's periodic returns. Thus, the only risk that matters to investors in an asset can be quantified by that asset's covariance with the market portfolio.⁹

Step 4: Asset Pricing and Expected Returns. As all investors hold their entire wealth in the market portfolio (plus or minus a riskless asset), the variance in the market portfolio quantifies the risk that all investors are exposed to. Therefore, the risk premium in the expected return to the market portfolio, divided by the variance in the return to the market portfolio, quantifies the "price of risk," that is, the market's required expected return risk premium *per unit of risk* (variance). Multiplying this price of risk by the amount of risk in each asset (that is, the risk that matters to investors, the asset's covariance with the market) gives the market's required risk premium in each asset's expected return:

$$E[r_i] = r_f + RP_i = r_f + COV_{iM} \left(\frac{E[r_M] - r_f}{S_M^2} \right) \quad (1a)$$

where $E[r_i]$ is the market's equilibrium required expected return to asset class i (constituent of the market portfolio), $E[r_M]$ is the market's equilibrium required expected return to the market portfolio as a whole, r_f is the risk-free interest rate, COV_{iM} is the covariance between asset class i and the market portfolio, and S_M^2 is the variance (volatility squared) of the market portfolio.

In just four steps, we have arrived at Formula (1a), the Sharpe-Lintner CAPM. We have simplified a bit in order to present a less technical and hopefully more intuitive derivation of this Nobel Prize-winning theory. But this is the complete model. It is important to understand where it comes from because such understanding can enable you to recognize both how it *can* be useful and the circumstances in which it *cannot* be validly used.

⁸ A simplified way of presenting essentially this same step in the model's derivation is to posit the concept that the entire capital market can be represented by a single "representative investor" who holds the market portfolio and who determines the pricing in the market.

⁹ Recall that the variance of the portfolio is the sum of all the pairwise weighted covariances among the portfolio's constituents:

Thus, the marginal contribution of any one constituent to the overall portfolio variance, per unit of weight of the constituent in the portfolio is:

$$VAR_p = \sum_{i=1}^N \sum_{j=1}^N w_i w_j COV_{ij}$$

The Sharpe-Lintner CAPM is usually expressed in a slightly different form than formula (1a). The more familiar formula uses a more intuitive measure of risk called **beta**. Beta simply normalizes the asset's risk as a fraction of the overall average or market risk (variance in the market):

$$\text{Beta}_{iM} = \frac{\text{COV}_{iM}}{S_M^2}$$

Using beta, the CAPM as expressed in formula (1a) becomes

$$E[r_f] = r_f + RP_i = r_f + \text{Beta}_{iM} (E[r_M] - r_f) \quad (2a)$$

In words, the CAPM says that an asset's expected return risk premium is directly proportional to its beta. Any asset's risk premium equals its beta times **the market price of risk**, which is observable as the expected return risk premium on the market portfolio, that is, the expected return on all investors' overall wealth over and above the risk-free interest rate.

As indicated in Exhibit 23-3, the CAPM provides financial economists and investment practitioners with a rigorous foundation for, and a way to formalize and quantify, the **security market line (SML)**, the relationship between risk and return that we first introduced way back in Chapter 3 where we called it "financial economics in a nutshell."

It is interesting to note that a mathematically equivalent way to express beta is as the correlation between asset class i and the market, multiplied by the ratio of asset i 's volatility divided by the market portfolio's volatility:

$$\text{Beta}_{iM} = \frac{S_i}{S_M} C_{iM} \quad (2b)$$

This makes it clear that the risk premium in an asset's required expected return is directly proportional not only to its risk as measured by the asset's own volatility (S_i), but also to its correlation with the market wealth portfolio (C_{iM}). An asset can be quite volatile, yet it still may not provide or require a high expected return if it is not highly correlated with the market. On the other hand, even

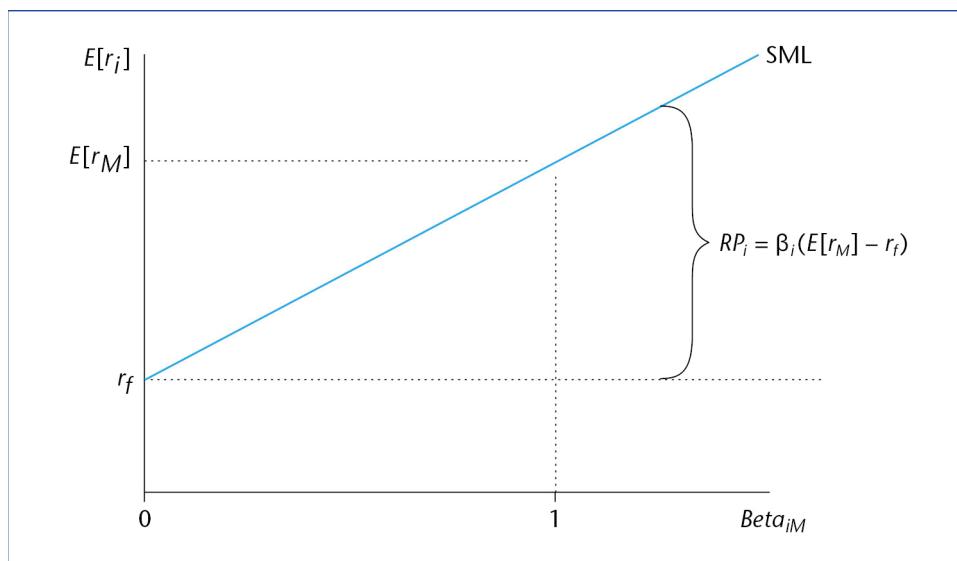


EXHIBIT 23-3 The CAPM in Graphical Form (the Security Market Line).

an asset with relatively low volatility may command a relatively high-risk premium if it is highly correlated with the market because in that case, it does not provide investors with much diversification benefit in their wealth portfolios. Risk is defined by the *interaction* of the asset's own volatility and its correlation with the *market*, that is, with investors' wealth portfolios.

23.2.2 THE MAIN POINT IN THE BASIC CAPM

The main investment insight provided by the CAPM is the *irrelevance of, and therefore lack of compensation for, diversifiable risk*. The CAPM suggests that, as covariance with the market portfolio is the only risk that matters to the capital market, it is therefore the only risk that will be priced in equilibrium. Risk that cannot be diversified away is referred to in asset pricing theory as **systematic risk**. Diversifiable risk, that is, the component of an asset's own total variance in excess of its covariance with the market, is referred to as the asset's **specific** (or **idiosyncratic**) **risk**.

$$S_i^2 = COV_{iM} + VAR[\varepsilon_i] \quad (3)$$

where $VAR[\varepsilon_i]$ is the variance of asset i 's idiosyncratic return component.¹⁰ Since idiosyncratic risk can be diversified away, it does not matter to the investor.

Particularly after our discussion of portfolio theory in Chapter 22, you can probably see that this point about non-pricing of diversifiable risk is actually broader and deeper than the specific Sharpe-Lintner CAPM itself. No matter how the risk that matters to the capital market is defined, whether it is well measured by the traditional CAPM beta or not, it is only such risk that will be priced, by definition. Under the (rather reasonable) assumption that virtually all investors can diversify away risk that is very specific to any one asset or group of assets, it therefore follows that specific risk will not be priced.

An important investment strategy is implied by this result: investors should generally not hold undiversified portfolios. Since investors are not compensated by the market (in the form of expected return) for exposing themselves to specific risk, they should get rid of all such risk by diversifying their wealth portfolio. Indeed, strictly speaking, the Sharpe-Lintner CAPM says that you should not hold any portfolio other than the market portfolio (scaled down to your wealth, of course, and levered up or down to conform to your risk preferences).

23.2.3 ISN'T THE CAPM WRONG?

At this point, students often cry foul. “The CAPM can't be right,” they say, “it's based on false assumptions.” Indeed it is. There is no such thing as a truly riskless asset. Markets are not perfectly efficient. Not all investors share the same expectations about the future. And we do not all hold the same portfolios (that is, relative risky asset weights in our wealth investments). Thus, *obviously*, the CAPM is “wrong” in the sense that it is not a *complete* description of reality of equilibrium asset pricing. Good empirical tests should be able to detect this fact, and indeed they do.¹¹

The doubling (multiplication by two) drops out in the comparison of the *relative* amount of risk contributed by each constituent. So the relative risk that matters is completely quantified by the covariance with the market.

$$\partial VAR_p / \partial w_i = 2 \sum_{i=1}^N w_j COV_{ij} = 2 COV_{ip}$$

¹⁰ By definition, idiosyncratic risk is the component of the asset's volatility that is not correlated with anything else. In a time-series statistical regression of asset i 's periodic returns onto the market portfolio's periodic returns, $VAR[\varepsilon_i]$ would be the variance of the *residuals* from the regression (sometimes referred to as the errors because they are the difference between the returns predicted by the regression and the actual returns of asset i).

¹¹ In a strict sense, the CAPM has been “flunking” empirical tests at least since the early 1970s, when famous studies by Black, Jensen, and Scholes (Black, F., M. Jensen, and M. Scholes. 1972. “The capital asset pricing model: Some empirical

In an important sense, however, this criticism misses the point, not only about the Sharpe-Lintner CAPM, but also about the fundamental nature and purpose of equilibrium asset price modeling in general, and indeed, about the purpose of virtually all scientific theory (if we may wax poetic). By this criterion of “wrongness,” Newtonian physics is wrong, as it is built on simplified (false) assumptions that ignore the relativity of time and space. Yet who can deny that we learn a lot that is quite useful from the Newtonian model, both at a practical level and at a deeper, more fundamental level. Virtually all industrial and engineering advances up to at least the middle of the twentieth century were based purely on Newtonian physics. Like Newtonian physics, the CAPM continues to be used widely in practice, not only by academics, but also by professionals in the investment industry.

What the CAPM loses as a result of its unrealistic assumptions, namely, the ability to model the world completely, is more than made up for by what the model gains by these assumptions, namely, the ability to simplify the world so that we can understand it better. The elegance in the basic CAPM enables us to obtain insights that we could not from a more complex model.

So, let’s consider the value and usefulness, as well as the limits and shortcomings, of the CAPM from a more sophisticated perspective. In so doing, we can begin to see how it may be modified or extended particularly to improve its relevance for real estate.

23.2.4 STRENGTHS AND WEAKNESSES IN THE BASIC CAPM

The CAPM, like any valuation model, can be viewed from either a normative or descriptive perspective. From the normative perspective, the model is addressing what *should* be, under the given assumptions. From this perspective we can gain what might be called wisdom, insight that improves our understanding. For example, the general point about the non-pricing of diversifiable risk is a key insight. The extreme suggestion that all investors should hold the same portfolio no doubt goes too far. But the suggestion that all investors should diversify their wealth is a good prescription and holds widely in reality. Even modest investors can hold broad market indices using ETFs (exchange traded funds) like SPDR comprised of all stocks in the S&P 500.

From a descriptive perspective, the CAPM is addressing what *is*, out there in the real world. From this perspective, we hope to gain a practical tool to predict what will happen, a model of cause and effect in the real world. In this respect, the CAPM, while not perfect, is useful in practice. Empirically, beta is not the whole story for explaining expected returns, but it is a big part of the story.

The CAPM is useful in practice presumably because the model’s assumptions, while simplifications of reality, are not terribly far from the truth. Although asset markets are not perfectly efficient, they are reasonably efficient most of the time (especially securities markets such as the stock market). Although all investors do not hold the same expectations or the same portfolios, most investors hold fairly similar expectations and fairly similar portfolios (especially large institutional investors).¹² As a result, the CAPM works well enough to be useful, at least at a broad-brush level, or with some fairly simple enhancements. Indeed, beta has become virtually a household word on Wall Street and is regularly estimated and reported as a matter of course by stock analysts. It is especially used in the equity mutual fund business.

On the other hand, the hopes that were originally held for the basic Sharpe-Lintner CAPM have not been completely fulfilled. We now recognize that a substantial part of the cross-section of expected returns across assets within the stock market is not well explained by beta. Other risk

tests”, in M. Jensen (ed.), *Studies in the Theory of Capital Markets*, New York: Praeger.) and Fama and MacBeth (Fama, E., and J. MacBeth. 1973. “Risk, Return, and Equilibrium: Empirical Tests,” *Journal of Political Economy* 81: 607–636.) were published. Equally influential were later studies published by Fama and French (Fama, E., and K. French. 1992. “The Cross-Section of Expected Stock Returns.” *Journal of Finance* 47(2): 427–465.) or Black (Black, F. “Beta and Return,” *Journal of Portfolio Management* 20: 8–18.).

¹² The development and tremendous growth of passive “index funds” in the mutual fund industry was originally stimulated and largely motivated by the efficient market hypothesis and the type of reasoning that is reflected in the CAPM.

factors, reflected in such indicators as the size of a firm and its book/market value ratio (the so-called “Fama-French factors”), have a strong influence on expected returns apart from a stock’s beta with respect to the stock market.

Nevertheless, it is important to keep in mind several caveats before excessively discounting the value of the basic CAPM. First, although beta (with respect to the stock market) is not the whole story about risk premiums within the stock market, it is an important part of the story. For example, once one controls for such other factors as firm size and book/market value ratio, beta describes a large portion of the remaining variation in average ex post returns to portfolios of stocks. Second, the additional risk factors that matter to the market apart from beta (e.g., the risk factors captured in such characteristics as firm size and current yield) may not *negate*, as much as they *refine*, the theoretical underpinnings of the basic CAPM.

For example, small firms and firms with high current yields or high book/market value ratios are often most susceptible to catastrophic failure in the event of macroeconomic downturns or capital market crises.¹³ Exposure to this kind of risk may not be well measured by firms’ covariances with respect to the stock market.¹⁴ Yet it is difficult for investors to diversify away such risk exposure because the downside outcomes are correlated with the overall economy and the capital market broadly defined. Investors demand an additional risk premium for this type of risk exposure, for fundamentally the same reason that they demand a risk premium for beta, namely, the difficulty of diversifying away risk in their wealth or welfare. Thus, we see that in its more fundamental implications, such as the non-pricing of diversifiable risk, the basic CAPM remains largely intact.¹⁵

23.3 APPLYING THE CAPM TO THE PRIVATE REAL ESTATE ASSET CLASS AS A WHOLE

The CAPM has relevant application to real estate investment in at least two different contexts or foci. Perhaps the most obvious application is to REITs and other real estate firms within the stock market. REITs will be addressed in detail in Chapter 26, but as regards the applicability of the CAPM to REITs within the stock market (with risk defined with respect to stock market risk factors), a couple of points should be noted. First, REITs tend to be low-beta stocks. Second, the CAPM seems to predict REIT returns in a comparable manner, and as usefully, as it is applied to most other sectors of the stock market in general. Since many REITs are small-cap or high book/market stocks, one

¹³ It is easier for a small firm to go completely out of business, in part because small firms may have less bargaining power with their creditors. Also, among the population of firms listed on the stock exchange, small firms and firms with relatively high book/market value ratios will tend more often than other types of firms to be distressed in some manner, or recovering from recent distress. This is because, in general, the stock exchange tends to attract firms that are relatively large and typically have a growth-oriented business strategy. The smaller-cap, lower-multiple stocks on the exchange therefore tend to include a larger proportion of “fallen angels.” Such firms tend to be more vulnerable to crises or negative shocks in the real economy or the financial markets than are larger, more highly priced firms.

¹⁴ In part, this may be because such risk may not be “normal” (that is, represented by a Gaussian probability distribution) in a statistical sense, as it may be due to large, discrete events rather than continuous incremental change or news. In a sense, this type of risk may be more akin to event risk or default risk in corporate bonds or mortgages than to beta or covariance-based measures of risk. Default risk is typically measured by multiperiod average return measures such as yield degradation (either ex ante or ex post), rather than by periodic HPR-based risk measures such as beta.

¹⁵ See Fama and French (1992) or (2004) for more depth of discussion (Fama, E. F., and K. R. French. 2004. “The Capital Asset Pricing Model: Theory and Evidence,” *Journal of Economic Perspectives* 18: 25–46.). We will also return to this issue in Section 23.4. It should be noted that it is also possible that investors “overreact” in some irrational or “behavioral” manner to some characteristics of some investments, such as small firm size and high book/market ratio, causing such investments to be chronically undervalued and hence to provide higher returns than what would be rationally warranted. Shiller’s “excess volatility” that was first noted in his 1981 article that we previously cited began a line of inquiry about so-called “behavioral,” as distinct from “rational,” models of the stock market. In this case, the classical CAPM (without the Fama-French extension) may represent a sort of normative model of what “should” be in a rational world rather than what actually is in the real world.

implication is that REITs tend to command relatively high expected returns even though their betas tend to be relatively low.

The more challenging and unique application of the CAPM to real estate, however, is to direct investment in the private property market. Even within this focus on privately traded real estate, it is useful to distinguish two further levels at which one might try to apply the CAPM. At a more specific level, one could try to apply the CAPM *within* the private real estate asset class, for example, to quantify risk and expected return differences between different types of property in different locations. Application at this “sectoral” level will be the subject of the next section in this chapter. But at a more macro-level, one could try to apply the CAPM broadly to the mixed-asset portfolio as a whole, encompassing private real estate as one of the major asset classes. It is to this macro-level application that we turn our attention in the present section.

In fact, there is some evidence in support of the basic Sharpe-Lintner single-factor CAPM in application to private real estate at a broad-brush level across asset classes. But we must be more careful in how we define and measure the market portfolio and how we compute real estate periodic return statistics. Consider the market portfolio first.

According to the CAPM theory, as we described its derivation in Section 23.2.1, the market portfolio should represent all the wealth of *all* investors, in other words, the aggregate wealth in the economy as a whole. Yet the CAPM is traditionally applied using the stock market alone as the “proxy” for this overall wealth portfolio or risk benchmark. Clearly, the stock market alone is not adequate as a market proxy when we are trying to quantify the risk in an asset class that is not even included in the stock market, such as private real estate.

The other problem with early attempts to apply the CAPM to real estate had to do with the type of data being used to calculate the real estate beta. As noted in Chapter 3, periodic returns time-series data are necessary to calculate co-movement statistics such as correlations and beta. Such data are readily available for individual stocks and for the stock market as a whole, based on highly liquid market values reflected in, for example, daily closing prices of stocks. In the case of privately traded real estate, however, the traditional periodic returns data were based on appraised values of properties, such as the NCREIF Property Index (NPI) in the United States.

Appraisals are estimates of property values typically derived by appraisers who are trained to “look backwards” in time to find prices of “comparable properties” that were sold in the past. Furthermore, indices or portfolios aggregating the appraised values of many individual properties may include properties that are not reappraised every period. Thus, many valuations in the index or portfolio may be “stale.”

The result is that, compared with a more liquid contemporaneous market value index such as the stock market, real estate periodic return indices tend to be *lagged* and *smoothed* across time. This results in a tendency to underestimate the covariance between real estate and indices such as the stock market.

Suppose we correct both of these problems, that is, the data problem and the market portfolio definition problem, and then apply the basic single factor CAPM at the broad level of the three major risky asset classes that we considered in Chapter 22: stocks, bonds, and real estate. In particular, let us suppose that an overall national wealth portfolio (NWP) consisting of equal one-third shares of stocks, bonds, and real estate would well approximate the theoretical market portfolio that the CAPM requires in principle.¹⁶

Now to compute our asset class risk and return measures according to the CAPM, let’s use the example risk and return expectations we employed in Chapter 22. These expectations are repeated in Exhibit 23-4.

Although these expectations are only example numbers, they are not untypical realistic expectations in terms of what most matters here, which is the relative risk premia across the asset

¹⁶ As described in Chapter 1 (Exhibit 1-4), the complete “pie” of the investible wealth in the United States is probably characterized by weights not too different from this assumption. The exact weights in Exhibit 1-4, aggregated within the three risky asset classes only, are 24 percent stocks, 21 percent bonds, and 10 percent commercial real estate.

EXHIBIT 23-4

Typical Risk and Return Expectations

	Stocks	Bonds	Real Estate
Expected Return ($E[r]$)	10.00%	6.00%	7.00%
Volatility	15.00%	8.00%	10.00%
Correlation with:			
Stocks	100.00%	30.00%	25.00%
Bonds		100.00%	15.00%
Real Estate			100.00%

classes. Furthermore, the real estate second moments in these expectations (real estate's volatility and its correlations with the other asset classes) attempt to take account of, and correct for, the type of data-smoothing problems we described before.¹⁷

We can use the expectations in Exhibit 23-4 to compute the covariances between each asset class and the NWP that we are using as the CAPM market portfolio. For example, the covariance between real estate and the NWP is 0.004983, computed as follows:¹⁸

$$\begin{aligned}
 COV[r_{RE}, r_m] &= COV\left[r_{RE}, \frac{1}{3}r_{ST} + \frac{1}{3}r_{BN} + \frac{1}{3}e_{RE}\right] \\
 &= \frac{1}{3}COV[r_{RE}, r_{ST}] + \frac{1}{3}COV[r_{RE}, r_{BN}] + \frac{1}{3}COV[r_{RE}, e_{RE}] \\
 &= \frac{1}{3}S_{RE}S_{ST}C_{RE,ST} + \frac{1}{3}S_{RE}S_{BN}C_{RE,BN} + \frac{1}{3}S_{RE}^2 \\
 &= \frac{1}{3}(0.10)(0.15)(0.25) + \frac{1}{3}(0.10)(0.08)(0.15) + \frac{1}{3}(0.10)^2 \\
 &= 0.004983
 \end{aligned}$$

Similarly, the covariances with the market portfolio are 0.009950 and 0.003733 for stocks and bonds, respectively. The variance of our equally weighted market portfolio is $S_M^2 = 0.006222$.¹⁹ Thus, we can compute the beta of each of our three asset classes with respect to the overall mixed-asset market portfolio as follows:

$$\begin{aligned}
 Beta_{ST} &= 0.009950/0.006222 = 1.60 \\
 Beta_{BN} &= 0.003733/0.006222 = 0.60 \\
 Beta_{RE} &= 0.004983/0.006222 = 0.80
 \end{aligned}$$

Now we see that the real estate beta appears to be rather reasonable. With a beta with respect to the NWP of 0.80 compared to the stock market's NWP beta of 1.60, real estate has half the risk of the

¹⁷ In addition to the smoothing issue *per se*, these expectations probably also reflect a somewhat longer-term perspective than one might get just by looking at annual frequency historical periodic returns data. Historical evidence suggests that over short intervals (such as quarters or years) there is less positive correlation (and perhaps more negative correlation) than what is presented here between real estate and the other asset classes.

¹⁸ Some basic algebraic properties of the covariance statistic are useful to keep in mind: $COV[aX, bY] = abCOV[X, Y]$, and $COV[X, (Y + Z)] = COV[X, Y] + COV[X, Z]$, where a and b are constants and X , Y , and Z are random variables.

¹⁹ Computed using formula (A.6) from the Chapter 22 appendix: $(1/9)[0.0225 + 0.0064 + 0.01 + 2(0.0036 + 0.00375 + 0.0012)] = 0.006222$.

stock market, according to the CAPM, and long-term bonds have just a little bit less risk than real estate. These betas are pretty much in line with the return expectations shown in Exhibit 23-4, which are broadly consistent with typical expectations in the real world.

To clarify this point, note that, by the expectations in Exhibit 23-4, the market portfolio has an expected return of 7.67 percent.²⁰ Now recall that these return expectations are based on an assumed risk-free rate of 3 percent. Plugging the betas into the basic CAPM gives the following expected returns for each of the three major asset classes:

$$\begin{aligned} \text{Stocks: } & 3.00\% + 1.60(7.67\% - 3.00\%) = 10.46\% \\ \text{Bonds: } & 3.00\% + 0.60(7.67\% - 3.00\%) = 5.80\% \\ \text{Real Estate: } & 3.00\% + 0.80(7.67\% - 3.00\%) = 6.74\% \end{aligned}$$

As seen graphically in Exhibit 23-5, the basic theoretical CAPM fits pretty well these risk and return expectations. The real estate expected return of 7 percent is just slightly above the CAPM security market line prediction of 6.74 percent (perhaps reflecting the need for an extra illiquidity premium compared with securities). The bonds' expected return of 6 percent is also slightly above the CAPM prediction of 5.80 percent. The stock market's expected return of 10 percent is slightly below the predicted SML. While the expectations we are using here are just example numbers, they are broadly reasonable and in line with typical real-world perceptions relative to a risk-free rate of 3 percent.

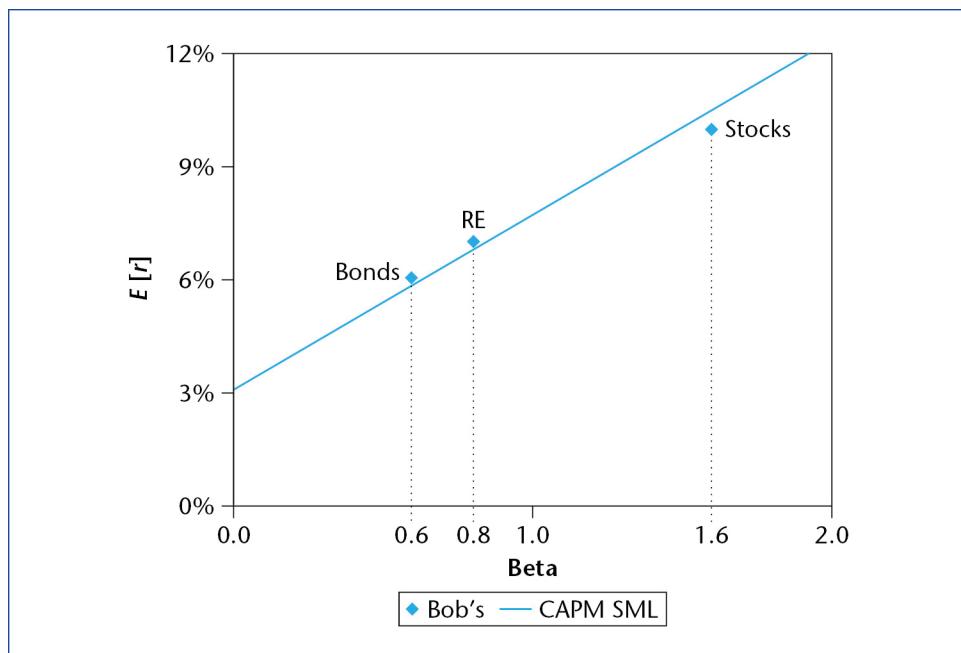


EXHIBIT 23-5 Typical Illustrative Expectations and the CAPM Prediction.

²⁰ This is computed as $(1/3)10\% + (1/3)6\% + (1/3)7\% = 7.67\%$.

In summary, the basic single-beta Sharpe-Lintner CAPM arguably works, in essence, at a broad-brush level *across* the asset classes. This is a level that is useful for broad strategic and tactical investment policy making for managers responsible for mixed-asset portfolios, that is, portfolios that potentially include all the major asset classes. It is also a level that is relevant for adjusting for risk the investment performance of managers or portfolios that encompass more than one broad asset class, for example, a real estate investment manager whose job includes allocation between direct private market investment and REIT equity investment.²¹

23.4 ATTEMPTING TO QUANTIFY RISK AND RETURN WITHIN THE PRIVATE REAL ESTATE ASSET CLASS

Can we apply the CAPM beyond the broad-brush asset class level to quantify risk and model the market's *ex ante* risk premiums at a more narrow and specific level within the private property market, for example, to quantify differences in the risk and return expectations for different types of property in different geographic locations? In other words, assuming *rational expectations* (that investors' *ex-ante* return expectations tend to actually manifest themselves in *ex-post* investment performance *on average and in the long run*), can the CAPM, or perhaps a similar but slightly enhanced equilibrium asset pricing model, explain the cross section of the long-run average returns across sectors or market segments within the overall private real estate asset class?

In general, it seems to be more difficult to apply the basic CAPM *within* asset classes than *across* them. This is true not only for real estate but also for the stock market. The classical single-beta CAPM does not explain much of the variation in *ex-post* returns across portfolios of individual stocks within the market.

A dramatic picture of this is seen in Exhibit 23-6, which reproduces a famous graph from an article by Eugene Fama and Kenneth French (2004). The exhibit in Exhibit 23-6 differs markedly from Exhibit 23-5. Each point in Exhibit 23-6 represents the long-run (1963–2003) average return and beta (with respect to the stock market) for one of ten portfolios defined by their average book/market value ratios. If the simple single-factor CAPM held, then all the portfolios should theoretically plot along the straight upward-sloping line shown in the exhibit. In fact, the portfolio risk/return points do not bear any relationship to the CAPM security market line. Instead, they plot very nicely according to their book/market value ratios (indicated by the portfolio label numbers), with portfolios that have lower market values (relative to the book or historical cost values of their net assets) tending to earn higher returns. So-called “value stocks,” the type that are preferred by investors of the “value style,” which typically have higher book/market ratios, provided higher returns during the 1963–2003 period, no matter what their beta was. In other words, “growth stocks” (typically having lower book/market ratios) were priced at a premium by the stock market, providing a persistently lower return (again, no matter what their beta).

Fama and French thus expanded the traditional CAPM to include in addition to the market beta two new “risk” factors as determinants of the expected return: the book/market ratio factor noted in Exhibit 23-6, and a “small stock” factor representing the “size” of the stock in terms of its market capitalization. In the case of the stock market, the empirical prominence of the Fama–French factors has led to the widespread acceptance and use of this “Three Factor Model.” The additional risk factors are referred to as “HML” (for high-minus-low book/market ratio portfolio returns) and “SMB” (for small-minus-big stock portfolio returns). It remains controversial among academic financial economists whether the Fama-French factors are important because of some sort of irrationality or “behavioral” phenomenon on the part of stock market investors, or whether the extra two factors in fact proxy for rational forms of risk perhaps as described in Section 23.2.4.

²¹ Specific quantitative methods for using the CAPM to adjust portfolio returns for risk are presented in all major graduate-level investments textbooks. See, for example, Bodie, Kane, and Marcus (2024), *Investments*, McGraw Hill, Ch. 24. A single-index model like the basic CAPM allows a Treynor-type measure to be employed for risk adjustment.

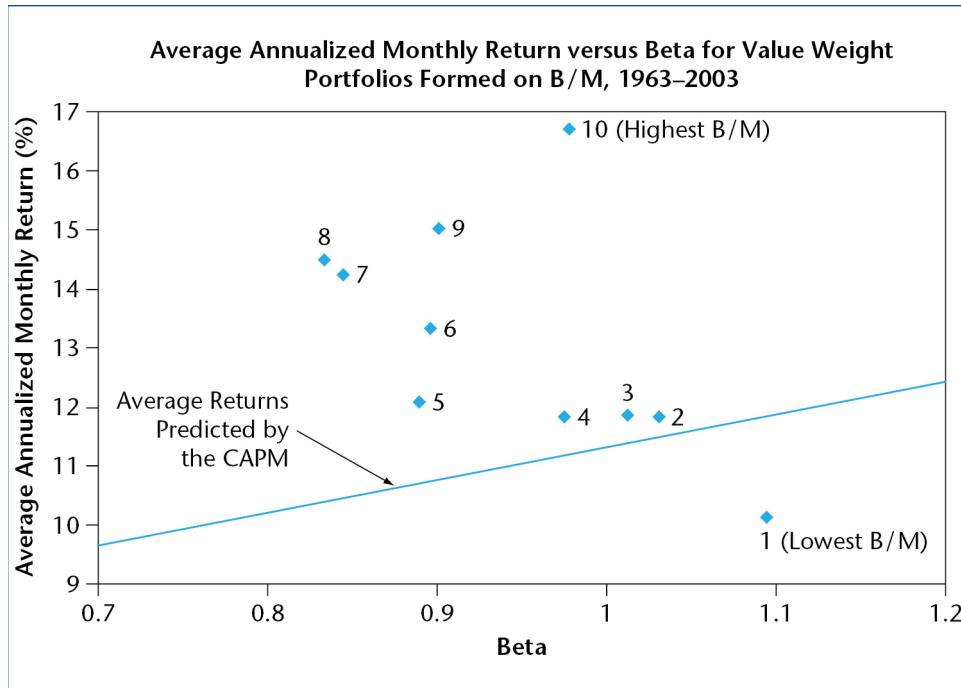


EXHIBIT 23-6 Stock Market Historical Fama-French Book/Market Value Portfolios Return versus Beta.

Source: Reproduced from Fama and French (2004), Exhibit 3.

The situation appears to be similar within the private real estate asset class. If anything, we are even less able to draw conclusions about causality and systematic patterns in real estate risk and return because there is much less historical data quantifying risk and return performance than in the stock market. As an example of the type of behavior one finds, Exhibit 23-7 summarizes some results from an examination of portfolios or subindices of NCREIF properties during 2000–2011.

Exhibit 23-7 shows the risk and return performance of 15 mutually exclusive and exhaustive subindices of the NCREIF Property Index (NPI). The 15 subindices are differentiated along two dimensions—property size (value) and usage type sector. The five sectors are CBD office properties, suburban office properties, industrial, retail, and apartment properties. The three breakouts by size have value cutoffs that vary by sector, so as to keep approximately one-third of all the NCREIF properties within each size category: small, medium, and big.²² The vertical axis plots the subindices' average quarterly ex-post total return risk premia (over T-bills) during 2000–2011 (45 quarters of data). The horizontal axis plots the subindices' betas with respect to the NPI, effectively their traditional CAPM type systematic risk with respect to the institutional investment property market represented by NCREIF (the large tax-exempt investment institutions in the United States, primarily pension and endowment funds, invested in real estate largely through specialized investment managers).²³

²² For example, within the CBD office sector, which tends to have the largest or most valuable properties, the “small” category consisted of properties below \$40,000,000 in appraised value and the “big” category was properties over \$110 million. At the other extreme, NCREIF industrial properties are much less pricey, and the “small” threshold was below \$7 million with the “big” threshold being above \$18 million.

²³ The smoothing and lagging present in the NPI has little impact on the betas of components of the NPI with respect to the NPI itself, because the smoothing and lagging is essentially the same on both sides of the regression and thus cancels

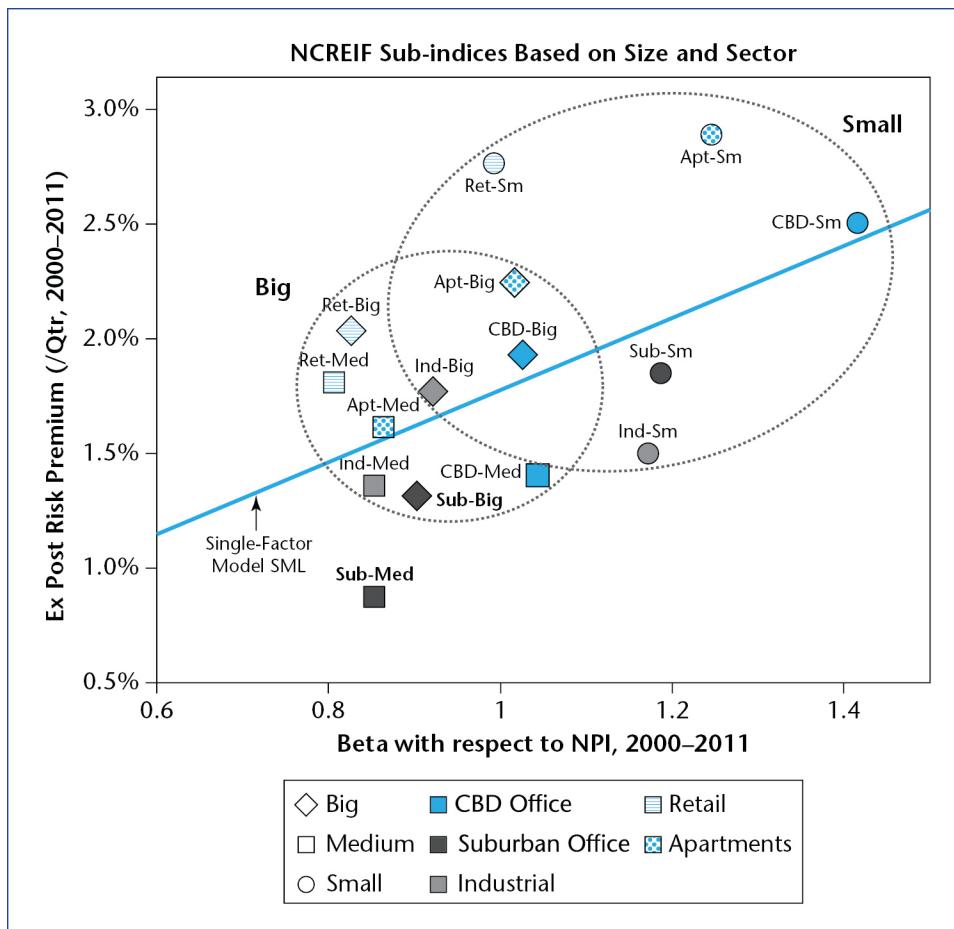


EXHIBIT 23-7 Risk and Return within Private Institutional Commercial Property, 2000–2011: NCREF Subindices Based on Size and Sector.

Source: Based on data from Jones, R. H. (2012). “Risk and return in institutional commercial real estate: A fresh look with new data.” Master’s Thesis, MIT Center for Real Estate.

If the classical single-factor CAPM well modeled the real estate risk and return in Exhibit 23-7 (and assuming that among NCREF properties the beta with respect to the NPI is generally proportional to the beta with respect to the overall national wealth portfolio as described in the previous section), then we should see all the points representing the subindices plot along a straight upward-sloping line similar to the single-factor Security Market Line (SML) indicated in the exhibit. In fact, the subindices’ risk/return performances are rather scattered, though not as perversely related to the CAPM as appeared to be the case with the stock market in Exhibit 23-6. In Exhibit 23-7, we do see some hint of a positive relationship between the average 2000–2011 returns and the betas. In fact, a statistical regression of the 15 risk premia onto their betas yields a strongly statistically significant positive coefficient on beta (with a value equal to 1.6 percent per quarter, reflecting the generally robust performance of commercial property investments over the 2000–2011 period). This provides some evidence in favor of the CAPM. However, the R-squared statistic of the regression

is only 27 percent, meaning that most of the variation in the subindices returns is not explained by their betas.²⁴

What if we try to expand the asset pricing model in the spirit of the Fama-French extra factors that seem to be so helpful in explaining the cross-section of stock returns?²⁵ In particular, let's add a size factor similar to the Fama-French SML factor, represented by the difference between the small property minus big property returns (each period). Another factor that institutional real estate investors claim is important is the metropolitan "tier" where the property is located. Top-tier cities disproportionately attract foreign capital, are relatively constrained in their physical supply elasticity, and may be more liquid for large investors such as NCREIF members. Labeled as *major* markets, during the 2000–2011 period measured in Exhibit 23-7, these were considered to include New York, Los Angeles, Chicago, San Francisco, Washington, D.C., and Boston. The industry then defines a *secondary* tier of cities that are also major targets for domestic institutional investors.²⁶ All other locations are categorized as *tertiary* markets for large institutional investors. The metro tier factor is thus constructed as the major minus tertiary returns (MMT). Finally, investors seem to be sensitive to the usage type sector directly, so we can define dummy variables representing the five property type sectors in Exhibit 23-7.

Now regress the 15 subindices' average 2000–2011 returns onto an expanded set of risk factors including their betas, their size and market tier loadings,²⁷ and property type sector dummy variables.²⁸ The result is that the R-square of the regression rises to over 90 percent,²⁹ with an insignificant intercept, so that the risk factors now explain most of the variation in returns. The market beta (with respect to the NPI) remains highly statistically significant, as are the betas with respect to both property size and metro tier, with positive signs (meaning that smaller properties, and "major" metros, have positive risk premia). In addition, at least some of the property type sector dummies are significant, most notably the apartment sector which commands a strong positive return premium.

We thus seem to have an asset pricing model that was rather effective at explaining the cross section of NCREIF property returns during 2000–2011. But there are several problems with drawing any major conclusions from such a study. For one thing, it is just one slice of the private real estate asset class. When we slice the NCREIF properties in different ways, for example, creating the subindices based on metro tier rather than on size, or when we go beyond the NCREIF population of properties to a broader representation of commercial investment properties such as those underlying the Moody's/RCA Commercial Property Price Index (CPPI), or if we look at an earlier historical period within the NCREIF Index, the model does not perform as well.

out in the estimation of the beta. Furthermore, it is plausible to assume that the betas of NPI components with respect to a broader market indicator such as the NWP described in the previous section would be proportional to their betas with respect to the NPI. Such proportionality is all we need for the CAPM to apply in theory.

²⁴ Of course, some of the dispersion of individual portfolios around the SML would always be expected, reflecting the realization of idiosyncratic risk in portfolios that are not perfectly diversified. But a risk factor that only explains a small portion of long-run average returns is not very useful.

²⁵ The study described in this paragraph is presented in Jones (2012), and is an extension of earlier work reported in Pai and Geltner (2007).

²⁶ For the study described here, applying to the 2000–2011 period, the "secondary" cities were defined as Atlanta, Miami, Dallas, Houston, Phoenix, Denver, San Diego, Seattle, Minneapolis, and Philadelphia.

²⁷ The term "loadings" here refers to the equivalent of "beta" only with respect to the extra factors returns instead of with respect to the market returns. For example, a subindex's size loading is determined by a longitudinal regression of the time series of the index returns (2000–2011) onto the time series of the factor returns (such as the SML returns during 2000–2011).

²⁸ For example, the apartment dummy variable equals one for the three (different size-based) subindices that are apartment properties, and zero for the other 12 subindices.

²⁹ Adjusted R-square over 85 percent.

A particular concern regards asset pricing factors that do not clearly make sense from the perspective of risk or other plausible concerns that investors might have. Two such factors stand out in the investment property performance results of recent decades. One is the metro tier factor in which the major markets out-perform the tertiary markets (MMT). If this reflects the rational pricing of a systematic risk factor, then it implies that investors regard the major markets as being more risky (or less liquid, or some other negative attribute from an investment perspective) than the tertiary markets. The other anomalous risk factor is the apartment effect. Do institutional investors really perceive large-scale apartment properties as more risky or less liquid than other types of property? Neither of these risk factors seems to make sense as such, and most investors would not claim them to be risk factors. This is borne out by surveys of investors' *stated* return expectations. Look again at Exhibit 11-8 back in Chapter 11, which shows the transaction cap rates from the MSCI data. Although cap rates are not quite the same thing as total returns (notably lacking the expected capital growth component), in general there is no reason to think that investors expect larger returns from apartments or properties in major markets, particularly among the larger "institutional" properties of the type in the NCREIF population.

When a risk/return analysis produces results that fit well empirically but do not make sense theoretically, one suspects an idiosyncratic historical result, a purely ex-post return result rather than the reflection of a persistent ex-ante pricing of risk. This is particularly true with short historical samples of data or narrow samples of properties to analyze, which is still what we are faced with in real estate. The 1990s and 2000s were favorable periods for apartments and major markets in terms of drawing institutional capital, and the result was a secular repricing of those assets, providing a windfall return to their investors. That does not necessarily mean that apartments or major markets are inherently or persistently viewed as more risky and therefore requiring risk premium in expected returns.

There are, however, two risk factors in the studies of private real estate risk and return whose effects on long-run average returns do make some theoretical sense. These are the market beta and the small property (SML) factor. Our presentation of the classical CAPM explains why beta makes sense as a risk factor for which investors should demand a return premium. In the case of smaller properties, these may be of concern to institutional investors because they are perceived as lacking liquidity, or perhaps because they are more difficult for large investment institutions to manage. Or smaller properties may command a premium due to poorer quality of information about them (a type of "uncertainty premium"). The small property return premium is perhaps the most persistent return premium within the private real estate asset class. It is observed in Exhibit 11-8 as the roughly 50-100 basis-point extra risk premium for "noninstitutional" property.

In summary, while there may be some explanatory power for some models of the cross-sectional variation in long-run asset returns within the private real estate market, in general results are weak or ambiguous, with the exception of the apparent premium for smaller (or "noninstitutional") properties.³⁰

³⁰ An interesting innovative perspective is taken in a study by Professor Liang Peng at the Pennsylvania State University. His study examines the property-level round-trip investment IRRs of thousands of NCREIF properties at the disaggregate level, instead of periodic returns of subindices or portfolios at the aggregate level. Peng's findings suggest strong relationships between some economic and financial factors and property-level returns. However, it is still not clear that such results reflect persistent ex-ante expectations or only ex post results reflecting the direct or indirect impact of the economic and financial variables (see Peng, L., 2016, "The Risk and Return of Commercial Real Estate: A Property Level Analysis," *Real Estate Economics* 44: 555-583).

TEXT BOX 23.1 IS IT RISK, OR IS IT EXPECTED RETURN? A QUESTION OF TERMINOLOGY AND INVESTMENT CULTURE

In the practice of macro-level real estate tactical investment analysis, the terminology can often be a bit confusing. In part, this confusion reflects real estate's unique investment "culture," as compared to traditional securities market investment practice. In particular, real estate investment analysts will often refer to differences in risk across sectors. For example, they may make statements such as: "The Atlanta apartment market seems particularly risky at present due to a current oversupply of new construction." It is important to distinguish risk as used in this context, from the concept of risk viewed from the perspective of mean-variance investment performance, the perspective that prevails in the stock market. The former use of the word *risk*, deriving from a space market analysis (e.g., the apartment market in Atlanta), actually means that the expected rent and occupancy in the relevant space market over the coming months or years is below "normal" in some sense (e.g., below the long-run equilibrium level). If (and it is a big if), property asset market prices in the subject sector have not responded to reflect this below-normal expectation in rent levels or rent growth (and assuming the expectation is valid given all the currently available information), then expected returns in the asset market sector will also be below normal.

But this does not necessarily imply anything about the risk in the future investment performance of the subject asset market sector. If we define risk from a mean-variance perspective, the way it is usually conceived in the capital market. From this latter perspective, risk refers to periodic return volatility (or derivatives of this volatility, such as beta, the risk that matters from an asset pricing perspective in terms of the investment market's required ex-ante return risk premium). For example, Atlanta apartment buildings may have a relatively safe return in the sense that it may be almost a sure bet that their returns will be below normal (or normal, if apartment property prices already reflect the bad news about future rental growth). From the perspective of capital market and investment terminology, Atlanta apartments may or may not be overpriced, but this is a question largely of their expected return, not of their risk (volatility or beta).³¹

In a sense, you can think of risk as defined in the former (space market) usage as being more akin to the capital market concept of default risk as it applies in the corporate bond or mortgage market. Risk in this context refers to an expected gap (in an ex-ante probabilistic sense) between the normal cash flow (akin to the contractual cash flow in a bond or mortgage) and the realized cash flow ex-post. If the asset price does not reflect this probabilistic expectation (as in the par value of the bond or the contractual principal of the mortgage), then the probabilistic expected rate of return on the investment is less than its normal rate (e.g., the stated yield or contractual yield in the case of the bond or mortgage). Thus, application of the space-market-based concept of risk in a real estate tactical investment context suggests that such investment is viewed largely from an underwriting perspective like that which typifies the primary (new issuance) market for bonds or mortgages. This underwriting culture contrasts with the more trading-oriented perspective that prevails in the stock market and the secondary market for bonds. This is another example of how real estate investments (in particular, direct investment in the private property market) fall into a unique category different from traditional mainstream securities investments, often blending characteristics of both stock and bond investments at both the primary and secondary market levels.

³¹ Of course, the volatility or uncertainty surrounding the future values of Atlanta apartments could conceivably be above normal (meaning, above its typical or historical level). But it is difficult to quantify property market sector volatility precisely and reliably over time. It is unlikely that the property market holds a very precise or stable idea about the volatility of Atlanta apartments in this regard, at least within the private (direct) property asset market.

23.5 CHAPTER SUMMARY

This chapter is built on the foundation of portfolio theory to introduce you to equilibrium asset pricing models, a key tool used in both strategic and tactical investment decision making at the macro-level, and even at the policy implementation level of investment management performance evaluation. We attempted not only to review, at an introductory level, the use of such models in the stock market (which includes REIT applications), but also to go into more depth regarding applications of macro-level asset pricing theory to privately traded real estate assets in the property market. Asset pricing theory faces some of its most intellectually exciting challenges and potential rewards in its application to privately held property. In this regard, we pointed out the importance of distinguishing two levels of application. The broad, mixed-asset level deals with pricing *across* asset classes or types of investment products and arenas. At this level, an elegant single-index CAPM-type model probably works pretty well. The more specific sector level deals with pricing *within* the private property market. At this level, the difference between institutional quality and noninstitutional properties looms large, but beyond that, the evidence for a rational model explaining the cross-sectional differences in long-run average returns based on equilibrium asset pricing is weak and ambiguous (though research continues, and with ever more and better historical return databases).

KEY TERMS

- Strategic investment decisions
- Tactical investment decisions
- Capital asset pricing model (CAPM)
- Market price of risk
- Market portfolio beta
- Security market line (SML)
- Systematic (nondiversifiable) risk
- Specific (idiosyncratic) risk
- Segmented markets
- Risk factors
- Institutional quality real estate

STUDY QUESTIONS

Conceptual Questions

- 23.1 What are equilibrium asset pricing models? Describe three practical uses for such models.
- 23.2 Which was developed first, the MPT or the CAPM? Why did the development occur in that order?
- 23.3 Why is beta referred to as a normalized measure of risk? In this normalization process, what (by definition) has a beta equal to unity?
- 23.4 In a nutshell, what is the main point of the CAPM? What is an important practical investment strategy implication of this point?
- 23.5 Describe two different levels, or foci, at which we might hope to apply the CAPM to real estate.
- 23.6 The CAPM was originally (and is still primarily only) applied within the stock market. Nevertheless, describe the two types of corrections or customizations to this narrow stock market application that allow the classical single-factor CAPM to be applied most elegantly to real estate at the broad-brush (mixed-asset) level. (Hint: The failure to consider these two points frustrated early attempts to apply the CAPM to real estate in the 1980s.)

23.7 Describe some capital market imperfections that render the CAPM a less than complete model of reality, especially for application to real estate.

23.8 How well does the single-factor CAPM seem to work within the institutional quality commercial property asset class? Discuss some possible reasons for your answer.

23.9 Consider a simple multifactor asset pricing model with only two factors, the market portfolio and unexpected inflation. What would you expect to be the *sign* of the price of risk of each of these factors (+ or -)? (Hint: Other things being equal, would you expect investors to prefer an asset that is positively correlated with unexpected inflation?)

23.10 If the risk-free interest rate is 5%, the market price of risk is 6%, and the beta is 0.5, then, according to the classical single-factor CAPM, what is the equilibrium expected total return for investment in the asset in question?

23.11 Fill in the first row of the following table (the expected returns) based on the single factor CAPM, assuming that the market portfolio has an expected return of 10% and consists of equal one-third shares of stocks, bonds, and real estate, and assuming that the risk-free interest rate is 5%.

CAPM-Based Risk and Return Expectations

	Stocks	Bonds	Real Estate
Expected Return ($E[r]$)			
Volatility	20.00%	10.00%	8.00%
Correlation with:			
Stocks	100.00%	60.00%	30.00%
Bonds		100.00%	-10.00%
Real Estate			100.00%

24 Data Challenges in Measuring Real Estate Periodic Returns

CHAPTER OUTLINE

- 24.1 The Basics: Macro-Level Property Valuation
 - 24.1.1 Transaction Noise and Appraisal Error
 - 24.1.2 A Simple Example of Temporal Lag Bias
 - 24.1.3 The Trade-Off between Noise and Lag
 - 24.1.4 The Difference between Aggregate and Disaggregate Valuation
 - 24.1.5 Summarizing Macro-Level Valuation Error
- 24.2 From Value Levels to Returns: The General Nature of Performance Measurement Errors
 - 24.2.1 The Pure Effect of Random Noise
 - 24.2.2 The Pure Effect of Temporal Lag
 - 24.2.3 Putting the Two Effects Together
 - 24.2.4 What about the Total Return?
- 24.3 What Is Truth? Lags and the Timeline of Price Discovery in Real Estate
 - 24.3.1 Multistep Real Estate Timeline of Information Incorporation into Observable Value
- 24.4 Innovations in Commercial Property Returns Indices
 - 24.4.1 Overview of Types of Indices
 - 24.4.2 Repeat-Sales Indices of Commercial Property Price Performance
- 24.5 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The major types of errors that tend to be present in empirical real estate periodic return data.
- The characteristic signs of such errors and the effects such errors have on analysis relevant to macro-level investment decision-making.
- The different ways in which it is conceptually possible to define real estate values and returns, the temporal relationship among these different definitions, and the relevance of different definitions for various types of practical decision problems.
- The new price and return indexing innovations for tracking commercial property, including “repeat-sales” indices based directly on investors’ actual experiences in the private property market.

In Chapter 3, we introduced and defined the concept of periodic returns or time series of holding period returns (HPRs). At frequent and crucial points throughout this book, we have used such returns for real estate and other asset classes in a variety of ways. For example, in Chapter 1, we compared the investment performance of the various asset classes, considering both risk and return. To make such comparisons, we used the HPR series. In Chapter 11, real estate periodic returns were important in our analysis of real estate cost of capital. They were also important in our discussion of leverage in

Chapter 12. In Chapters 22 and 23, periodic returns were the raw material for computing or estimating the expected returns, volatilities, and correlations necessary to apply modern portfolio theory and develop measures of risk and models of risk valuation and equilibrium asset pricing. Moreover, we will see in Chapter 25 that HPRs are fundamental to measuring and analyzing investment manager performance. In short, particularly (but not uniquely) at the macro-level, modern investment theory and practice absolutely *depend* on the basic data raw material represented by the time series of periodic returns.

The existence, reliability, and meaningfulness of such return data are taken for granted in the securities industry. For stocks and bonds, the history of such data goes back to the early decades of the twentieth century. The nature of public exchange trading of securities makes it easy to observe and measure periodic returns reliably and precisely, even at frequencies as high as daily for many securities. The existence of such vast quantities of high-quality data is a gold mine for the science of financial economics, the likes of which are the envy of all branches of social science.

In the 1970s, real estate investment industry leaders realized that for the investment establishment to perceive private real estate with a degree of credibility and legitimacy approaching that of stocks and bonds, it was necessary to compute, compile, and disseminate a series of real estate periodic returns. So, since the 1970s, periodic returns on commercial property have been reported and used in the investment industry. The “flagship” **NCREIF Property Index (NPI)** dates from the beginning of 1978, reporting quarterly total returns as well as income and appreciation components ever since, a history that now stretches for almost half a century.

At first, academics and industry traditionalists viewed the real estate returns series skeptically, and their use was not widespread. Over time, they have gradually become better constructed, more useful, understood, and widely accepted and used. The 2000s have seen an increase in the sources and types of commercial property returns data in the United States, as well as an extension of at least basic appraisal-based data in several other countries. As a result, in the United States, as of now, the three major types or sources of commercial property returns time series data, at least for the capital returns or asset price changes that are most important, are (1) appraisal-based, (2) transactions-based, and (3) stock market-based. Each type of data has different strengths and weaknesses, and some can be applied in some situations but not others. Thus, knowledge of all three of these data sources is an important part of the modern educated real estate investor’s toolkit, and in this chapter, we will introduce you to all three.

But first, it is important to recognize some of the unique issues with private real estate returns data, such as lag bias and statistical noise (and the trade-off between those two). Thus, the chapter will begin with the basics of commercial property returns measurement, including some consideration of valuation at both the disaggregate (individual property) level and the aggregate (collections of properties) level.

24.1 THE BASICS: MACRO-LEVEL PROPERTY VALUATION

Although this chapter is about indices or time-series data on the periodic returns to populations or portfolios of many properties in the direct (private) real estate asset market, the fundamental data from which most such return series are created are disaggregate, microlevel value indications of individual properties as of specific points in time. These individual valuations are aggregated or “averaged” to define a central tendency in the asset price change or return across the portfolio or population of properties or investors tracked by a given index. Therefore, before we consider aggregate returns, we must begin this chapter at a more fundamental level by building on concepts introduced first in Chapter 14, considering the valuation, or observation of value, of individual properties.

24.1.1 TRANSACTION NOISE AND APPRAISAL ERROR

In the supplemental materials of Chapter 14, we noted that observable empirical real estate values are “noisy”—that is, imprecise indications of “true value.” The primary empirical indication of value in the direct property market is, of course, the *transaction prices* of properties that actually

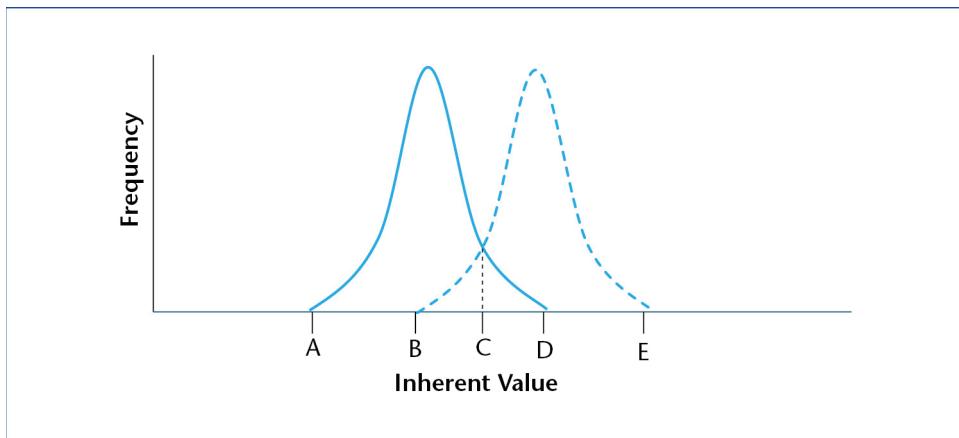


EXHIBIT 24-1 Buyer and Seller Populations, Reservation Price Frequency Distributions.¹

sell within the subject market or population of properties. But another important type of empirical value indication for properties is *appraised values* (or *appraisals*), estimates of specific properties' market values as of specified points in time made by professional real estate appraisers. However, appraisals are only *estimates* of value based more fundamentally on transaction price indications.¹

The key point made in Chapter 14A.3 is that both of these types of empirical value indications contain "error" relative to the "true" *market value* of the properties, at least if we view this process in a statistical sense for developing aggregate indices of market returns. We now need to add nuance to this point by noting that the error falls fundamentally into two major categories: purely *random error* (also known as *noise*) and *temporal lag bias*. Furthermore, when trying to design an optimal value estimation technique, there tends to be a natural trade-off between these two types of error. It is hard to reduce random error without increasing the lag bias, and it is hard to reduce the lag bias without increasing the random error.

To understand these fundamental points, recall the difference between market value and transaction price described in the online supplemental materials for Chapter 14. Exhibit 24-1 reproduces the key picture from Chapter 14A.3, showing overlapping reservation price distributions of potential buyers and sellers for some type or within some population of properties as of a given point in time.² The exhibit indicates that transactions in this market may occur in the overlap region where some buyers will have reservation prices at least as high as the reservation prices of some sellers, at prices ranging between values B and D. These prices will tend to be distributed around a value labeled C, which we described as the conceptual market value of the property type or the population of properties (e.g., adjusted for quality). Chapter 14A.3 describes how this value represents the classical economic concept of equilibrium value or market-clearing price. We can think of the value

¹ The role of appraisal and the relationship between appraisals and transaction prices varies in specific circumstances and also differs culturally from one country to another. The United States may represent a somewhat extreme case of market independence from appraisals: transaction prices reflect the market directly, while appraisals do so indirectly by filtering transaction price evidence. In other countries where the appraisal profession is strong (such as the United Kingdom), there may be more influence of appraisals on the transaction prices, with a dual direction of causality in which the professional appraisers ("valuers") play a more independent role in helping to determine market values. This could result in less systematic difference between appraised values and transaction prices. In some other countries, neither the appraisal profession nor the markets may be as well developed, and property exchanges may occur more rarely, with exchange prices governed by traditional formulaic value determinations. The primary professional organization governing the appraisal industry, which relates to most commercial and investment property in the United States, is The Appraisal Institute. In the United Kingdom and some other countries, the Royal Institution of Chartered Surveyors (RICS) is the major governing body.

² Recall from Chapter 14A.3 that the term "reservation price" refers to the private value at and above (below) that the potential seller (buyer) will stop searching or negotiating and agree to the transaction.

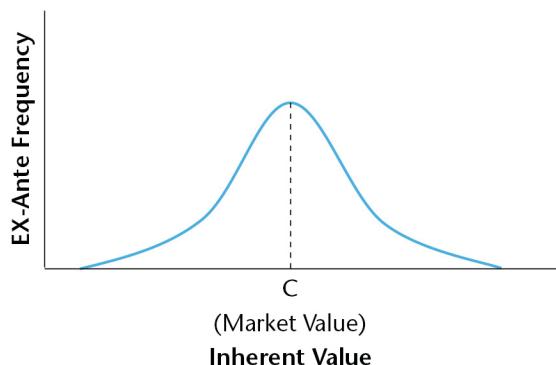


EXHIBIT 24-2 Theoretical Cross-Sectional Dispersion in Observable Value Indications.

C as representing the “true” value of a given point in time, with C being unobservable empirically. We can only observe valuations drawn from the probability distribution around C in Exhibit 24-2.

Suppose the empirical observations are actual transaction prices, the direct and fundamental indications of market value. In that case, the difference between any given price observation and the unobservable true market value is referred to as *transaction price noise* or *transaction price error*. By definition, this error will be unbiased (equally likely to be on the high side or the low side) as long as we hold time constant, that is, assuming all the transaction price observations occur at the same time.

Property appraisals are also dispersed cross-sectionally around true market values at any given point in time. If any two appraisers independently value the same property simultaneously, they will almost certainly not make the same estimate of market value. At least one of them must be “wrong” in the sense that their valuation differs from the true market value of the subject property. In fact, both appraisers are probably “wrong” in this sense. The difference between a given empirical appraised value and the (unobservable true) market value is called *appraisal error*. However, there is no implication that the appraiser has exhibited any incompetence, negligence, or impropriety.

Although appraised values are dispersed around the underlying true values, unlike transaction prices, the appraised value dispersion is not necessarily centered on the true value. In other words, appraised values may be biased at any given time. Such bias may result from very rational and proper professional practice on the appraiser’s part, given the nature of the empirical information available in the real estate market. The major bias likely to exist that concerns us in the present context is that appraised values tend to lag in time behind true contemporaneous market values. This is referred to as temporal lag bias.

24.1.2 A SIMPLE EXAMPLE OF TEMPORAL LAG BIAS

Suppose you are considering selling a property you own and hiring an appraiser to help you decide what price to sell it for. The appraiser offers you the following choice. She can give you a value estimate based on appraisal “Method A,” in which the value estimate has a 75 percent chance of being within 10 percent of the property’s true market value. However, Method A contains some **temporal lag bias** in that the expected (or most likely) value of the Method A appraisal is actually the true market value of the property six months ago. Thus, if true market values were, say, 3 percent lower six months ago, then Method A has a 3 percent bias on the low side. Alternatively, the

appraiser can base her appraisal on “Method C,” which will be unbiased regarding being completely up to date (i.e., the appraisal is equally likely to be above or below the current true market value).³ However, the Method C value estimate has only a 50 percent chance of being within 10 percent of the true market value of your property. Another consideration is that the more precise—but temporally biased—Method A will provide you with more solid historical evidence explicitly documenting your property’s estimated value. It will do this by providing you with more comparisons of your property with specific comparable sales transactions (known as “comps” in the appraisal field) that, in the appraiser’s judgment, are particularly comparable to your property. The appraiser says she will charge you the same price for either method. Which would you prefer?

The most typical answer is Method A, in part because its greater precision (75 percent instead of 50 percent chance of being within 10 percent of the true value) will probably be more useful to your decision problem relating to the *single* individual property that is the subject of the appraisal, and in part, because the greater specific documentation that method provides regarding your particular property can help you to persuade others of the value of the property. Because you control your own decision about the price at which you will actually agree to sell your property, you may be less worried about the bias in Method A. You probably have some sense of the direction of the temporal bias and maybe even a vague sense of its magnitude, just from your familiarity with how the market has been changing over the past six months. And, of course, you will try to sell the property for as high a price as possible, no matter what the appraiser says. In other words, the actual transaction price you end up with will reflect the current market more than the appraiser’s estimate.⁴

24.1.3 THE TRADE-OFF BETWEEN NOISE AND LAG

But now, let’s put ourselves in the shoes of the appraiser. Why did she have to offer you the choice between Method A and Method C? Why couldn’t the most precise appraisal also be the least biased (most up-to-date)?

The answer to this question lies in the fundamental relationship between statistics, time, and the generation of empirical data. Recall from the supplemental materials with Chapter 14 that we can make more accurate inferences about property market value, the more transaction price observations we have on which to base our estimate. Indeed, we said this reflects a fundamental law of statistical inference known as **the Square Root of n Rule**. The accuracy of our estimate will tend to be inversely proportional to the square root of the number of transaction-price indications of value that we can observe. If the appraiser can use six comps instead of just two, she can estimate the value that has $\frac{1}{\sqrt{3}}$ (or about 58 percent) as much error (because $6/2 = 3$). But comps are generated only when transactions actually occur, and transactions occur *over time*. In order to use six comps, the appraiser will have to reach three times as far back in time (on average).

The tradition of professional appraisal is to appraise individual properties using sales comps that are individually examined and considered for their relevance and appropriateness to the subject property. For an appraisal of large-scale commercial and investment properties, which tend to be much more

³ We’re labeling it “Method C” instead of “Method B,” because it is focused on the value “C” in the Exhibits, the unobservable but “true” market value.

⁴ Conceptually, this is true *by definition*. We define the “market value” of a given point in time as the mean of the distribution of potential transaction prices that could occur at that point in time. This does not imply that the transaction prices (and hence the market value) will not also be subject to some extent to the same type of partial adjustment or sluggish incorporation of new information that causes the appraiser to lag the market (perhaps rationally). We already noted in Chapter 14 (and raised the point again in Chapters 22 and 23) that real estate market values can have inertia and some degree of predictability due to a lack of perfect informational efficiency in the market. The difference in this regard between the market values and the appraisals is one of degree, and as noted previously, this degree may vary in different situations and culturally across different countries.

heterogeneous than houses and to sell much less frequently, this tradition effectively limits the comps sample typically to only a few properties sold in the local market area. The selected comps are supposed to be very similar to the subject, and their differences from the subject are scrutinized and carefully considered. This approach considers unique, idiosyncratic features of the subject property, eliminating the most important potential source of purely random error. However, the trade-off is that such comps are scarce in time, and the appraiser must, therefore, reach back relatively far in time.

Furthermore, appraisers tend to be conservative in adjusting for changes in market conditions, and rightly so. Such adjustments are subjective, and appraisers must objectively document their value estimations.⁵ In fact, statistical and cognitive theory both agree that it is rational to adjust one's prior opinions only partially in the face of new evidence that is less than definitive. If the appraiser truly ignored all past value indications and considered only transactions that have occurred, say, within the past month, no matter how far removed or different from the subject property these transactions were, and no matter how few such transactions were (e.g., maybe there is only one, or none!), then she would probably be considered negligent by the standards of the appraisal profession. Indeed, such an appraisal is likely not accurate for the specific subject property being appraised.

24.1.4 THE DIFFERENCE BETWEEN AGGREGATE AND DISAGGREGATE VALUATION

Now let us return to our example of the appraiser who offers you the choice between the “Method A” and “Method C” appraisals. Suppose the purpose of the appraisal is no longer to help you decide on the price of a single specific property but rather to help you obtain an *up-to-date* estimate of the *aggregate* value of a large portfolio of properties whose value is being tracked by an index, by combining this appraisal with similar appraisals done at the same time of all the other properties in the portfolio. Now, which method would you choose?

In this case, if you answered Method C, the less accurate one (for the individual property appraisal), you would be correct! This is because the purely random error in Method C will be largely diversified away at the portfolio or index level due to the rather powerful effect of the Square Root of n Rule in a large portfolio, where n is indeed a very large number. For example, if the portfolio contains 100 properties that are each being independently appraised each period, then the aggregate portfolio valuation relevant for index construction (made as the sum or arithmetic average across all of the individual property valuations) will contain only one-tenth ($1/\sqrt{100}$) the purely random error of the typical individual appraisals. By contrast, the temporal lag bias in Method A is systematic and will not diversify out of the portfolio valuation. The portfolio valuation will contain as much lag bias as the average individual property appraisal.

24.1.5 SUMMARIZING MACRO-LEVEL VALUATION ERROR

The preceding discussion of property valuation error can be summarized graphically in Exhibit 24-3. In the chart, the two axes represent the two types of error, arranged so that the farther out along the axes (away from the origin), the less the error (and the better the valuation). The horizontal axis represents greater precision in the value estimate, that is, less purely random error. Points farther to the right on the axis have less noise. For example, this dimension might be measured by the inverse of the standard deviation of the value estimates' purely random error component as a percent of the true property value. At least in principle, a point on this axis represents perfection, eliminating all random errors, indicated by the short vertical slash and the *0 percent* value label. The vertical axis represents greater “currentness,” or less temporal lag bias, the further up you go along the axis. Again, there is a point of theoretical perfection where zero lag bias exists in the value estimations, indicated by the *0-month* point on the vertical axis.

⁵ This may be becoming less of a problem with the advent of more and better transactions-based price indices for commercial property. (See discussion later in this chapter.)

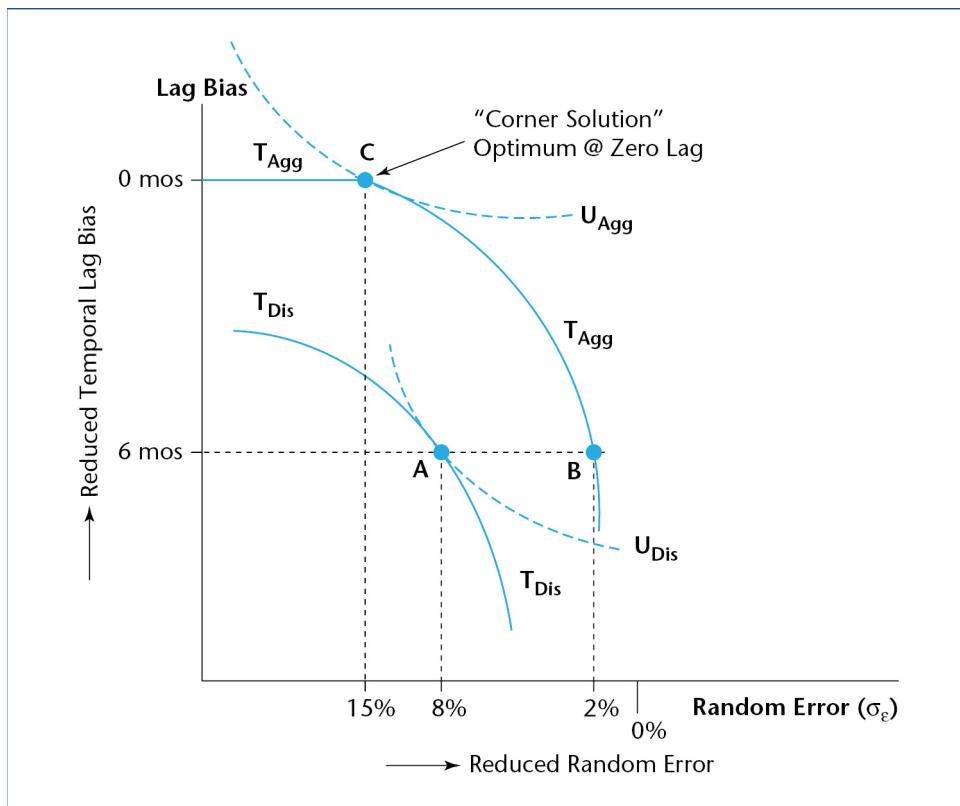


EXHIBIT 24-3 Noise-versus-Lag Trade-Off Frontiers with Disaggregate and Aggregate Valuation Methodologies.

The dashed convex curves are *isoquants*, or indifference curves of constant utility, from the users' perspective of the value estimation information. As users of property value estimates dislike both random error and lag bias, indifference curves that are farther up and to the right (farther "northeast" in the chart) represent higher levels of utility, like contour lines on a map of a mountain whose peak is in the northeast corner of the chart. The isoquants are convex (bending outward and upward so as not to intersect the axes) because of declining marginal utility for either sort of accuracy. Once a relatively high level of precision is obtained (low noise toward the right side of the horizontal axis), reducing any significant lag bias in the value estimate is more useful than adding another increment to the already high precision of the value estimate. Similarly, once the index is quite up to date (low lag bias, high up on the vertical axis), reducing any significant noise in the value estimate is more useful than reducing the lag bias by another increment.

The exact shape and slope of the utility isoquants will depend on the user and the use of the property-value estimation. For example, most users of individual (disaggregate) property appraisals place a premium on precision and don't care much about lag bias (such as the individual property appraisal to help advise a potential seller we described earlier). This would be reflected by a utility function with relatively steep isoquants tilted more clockwise or vertically, like the U_{Dis} curve in the exhibit. On the other hand, many uses of aggregate appraisal (such as constructing an index tracking a population of properties) may care more about avoiding lag bias. This would be reflected in a more shallow-sloped or horizontal isoquant like the U_{Agg} curve in the Exhibit.

The thick, solid, concave curves in Exhibit 24-3 represent the noise-versus-lag trade-off frontier, as provided essentially by the Square Root of n Rule. Points above and to the right of the trade-off frontier are not feasible and cannot be obtained. The exhibit shows two different value estimation

frontiers, representing two different property value estimation circumstances. The curve closer to the origin and farther to the left, labeled T_{Dis} , is the noise-versus-lag trade-off frontier that is relevant for traditional disaggregate, individual property appraisal, in which only a few comps are available to the appraiser to estimate the value of a single, specific property.

The kinked frontier farther up and to the right labeled T_{Agg} , including the straight horizontal section along the zero-lag boundary and the steeper curve approaching the zero-noise boundary, is the frontier that is relevant for macro-level index construction. Because the macro-level index covers many individual properties in the population or market, the Square Root of n Rule pushes the trade-off frontier far out to the right, near the zero-noise boundary. Even with no lag bias (only current transaction price information used in the value estimation), the size of the aggregate population of properties tracked by the index can give it a relatively low noise level. As it is impossible to have less than zero temporal lag bias, the aggregate trade-off frontier is kinked at the zero-lag boundary.

There is an important implication of the shapes and locations of the disaggregate and aggregate value estimation trade-off frontiers. The optimal balance between lag bias and random error is always found when the trade-off frontier is tangent to (that is, parallel to and just touching) one of the utility isoquants. This will be the feasible point that achieves the highest possible utility for the value estimation method. The optimal disaggregate-value-estimation method will likely involve some considerable degree of temporal lag bias in order to get the random error down to an acceptable level. This is indicated by point A , the point at which the T_{Dis} frontier is just tangent to the U_{Dis} level of valuation utility. This appraisal method produces some purely random errors and some lag bias. For example, consistent with our previous illustration (and probably not atypical of commercial property appraisal in the United States), optimal (Method A) type disaggregate appraisal might have six months' worth of average lag bias and 8 percent of purely random error (as indicated in the exhibit). Now, suppose a macro-level return index is produced simply by taking the sum or arithmetic average of such appraisals that are optimal at the disaggregate level. In that case, the result for the index will be a point like B , with less noise (due to the larger sample size inherent in the aggregation process) but the same amount of lag bias. But B is clearly not the optimal point for the aggregate index, at least from the perspective of the aggregate-focused users with utility function U_{Agg} . The optimal noise-versus-lag trade-off at the disaggregate level is different from the optimal noise-versus-lag trade-off at the aggregate level.

The optimal aggregate value estimation methodology will tend to occur at or near the “corner solution” with zero or very little lag bias at a point like C , which contains more noise than B (15 percent instead of 2 percent, which might be typical with current commercial property indices in the United States), but virtually zero lag bias. This reflects the difference in your choice of the appraisal methods discussed in Sections 24.1.2 and 24.1.4, wherein, in the latter case, you preferred the less precise appraisal Method C because it had less lag bias. In practice, an optimal aggregate valuation method such as point C in Exhibit 24-3 will usually be represented by some sort of regression-based transactions price index of the type that has only recently become available in the United States as described in Section 24.4.2 later in this chapter.

24.2 FROM VALUE LEVELS TO RETURNS: THE GENERAL NATURE OF PERFORMANCE MEASUREMENT ERRORS

Now that you understand the nature of the two types of errors that characterize property valuations, let's consider how these errors affect the main subject of this chapter, indices or time series of property-level periodic returns in the private real estate market.

To build your intuition in this matter, we will first consider the pure effect of **random noise**, then the pure effect of temporal lag bias, and finally, the effect of the two types of errors combined.

24.2.1 THE PURE EFFECT OF RANDOM NOISE

Suppose we know the “true” appreciation returns for a specific market. In reality, the “true” returns are unobservable. However, we *can* observe transactions in that same market. Furthermore, we can

ensure that any differences in composition in the transacted properties will not impact the returns, so we are comparing apples to apples (i.e., constant quality; see Section 24.4.2 on how we do this). Still, transaction “noise” will impact these observable appreciation returns. How will the empirically observable transaction-price-based annual appreciation returns differ from the unobservable appreciation returns to the same data based on the true market value? What does this effect of random noise look like?

The answer to this question is depicted visually in Exhibit 24-4. Panel 24-4A shows an example of the effect of noise on returns based on real-world data. Time is on the horizontal axis, with the periodic returns indicated on the vertical axis. The solid line in the chart is the true return realization over time. The dotted line is the empirically observable return realization. The Exhibit on the left is based on 500 observations, whereas the Exhibit on the right is based on 200 observations.⁶

Notice that the noisy direct price observation returns are considerably more volatile. Furthermore, the fewer observations, the more volatile the series gets. The empirical returns based on 500 observations have a standard deviation of 4.5 times larger than the “true” volatility of that same market. For the returns based on 200 observations, the volatility is 8 times as large. The empirical returns are also spikier or more saw-toothed in appearance, with large returns tending to be followed by opposite-signed returns. This is typical of excess volatility that reflects overshooting or random errors in the observed value levels that are subsequently corrected in later observations. This is seen more directly in Exhibit 24-4B, which shows the same history, only in the value levels rather than the returns. For example, the empirical valuation in one period might have a positive error (valuation higher than the true value), causing the valuation in the next period to likely at least partly correct that error, generating a negative return (value change) in the next period. The result would be a large apparent negative return due purely to the difference in errors when the true return was near zero. In general, note that random noise causes the observed value level series to sort of “dance” or “vibrate” around the unobservable true value level series. In the long run, both series should coalesce.

What is the effect of this type of return noise on the returns time-series statistics that are of most interest to investors? Well, it does not change the expected value of the periodic return because the expected value of the random valuation error is zero at both the beginning and end of each return reporting period. Hence, the expected value of the difference in errors across each period is zero, and it is this difference that exhibits in the periodic returns. However, random valuation error does change the volatility, that is, the standard deviation of the periodic returns across time. In particular, it increases the volatility, adding spurious “extra” volatility to the returns over time. Noise also affects the autocorrelation in the returns time series, that is, how the periodic returns are correlated across time.

In particular, noise reduces any positive first-order autocorrelation (that is, the correlation between consecutive returns) that might exist in the true returns and may make uncorrelated true returns appear to have negative autocorrelation.⁷ Noise does not affect the theoretical covariance between the index returns and any exogenous series because, by definition, a purely random variable has no covariance with any other series. Therefore, noise does not affect the theoretical beta (or systematic risk) as we defined it in Chapter 23.⁸ Finally, combining these **volatility** and **covariance effects**, noise reduces

⁶ The dashed line represents the returns according to a repeat sales model based on data for apartment sales in a specific census tract in Manhattan (source: REZITRADE). The repeat-sales model will be discussed below in Section 24.4.2. The solid line is the index after signal extraction. The method of signal extraction is discussed in Section 24.4.2 as well. For more details on this specific model and data, see Francke, M.K., L. Rolheiser, and A. Van de Minne, 2023. “Estimating Census Tract House Price Indexes: A New Spatial Dynamic Factor Approach,” *The Journal of Real Estate Finance and Economics*, 70(3): 483–514. Please note that we call the index the “true” index after signal extraction. In reality, the “true” index is unknown and unobservable. However, we do know that such signal extraction indices should be much closer to the true index than the noisy direct price observations indices.

⁷ The autocorrelation “signature” of pure random noise is negative 50 percent in the first-order autocorrelation of the returns (first differences of the value levels), and zero for higher-order autocorrelation.

⁸ However, noise does make it more difficult to precisely estimate the true beta using empirical returns data. In other words, it does not affect the expected value of the beta coefficient estimate in a regression, but it does increase the standard error in that estimate.

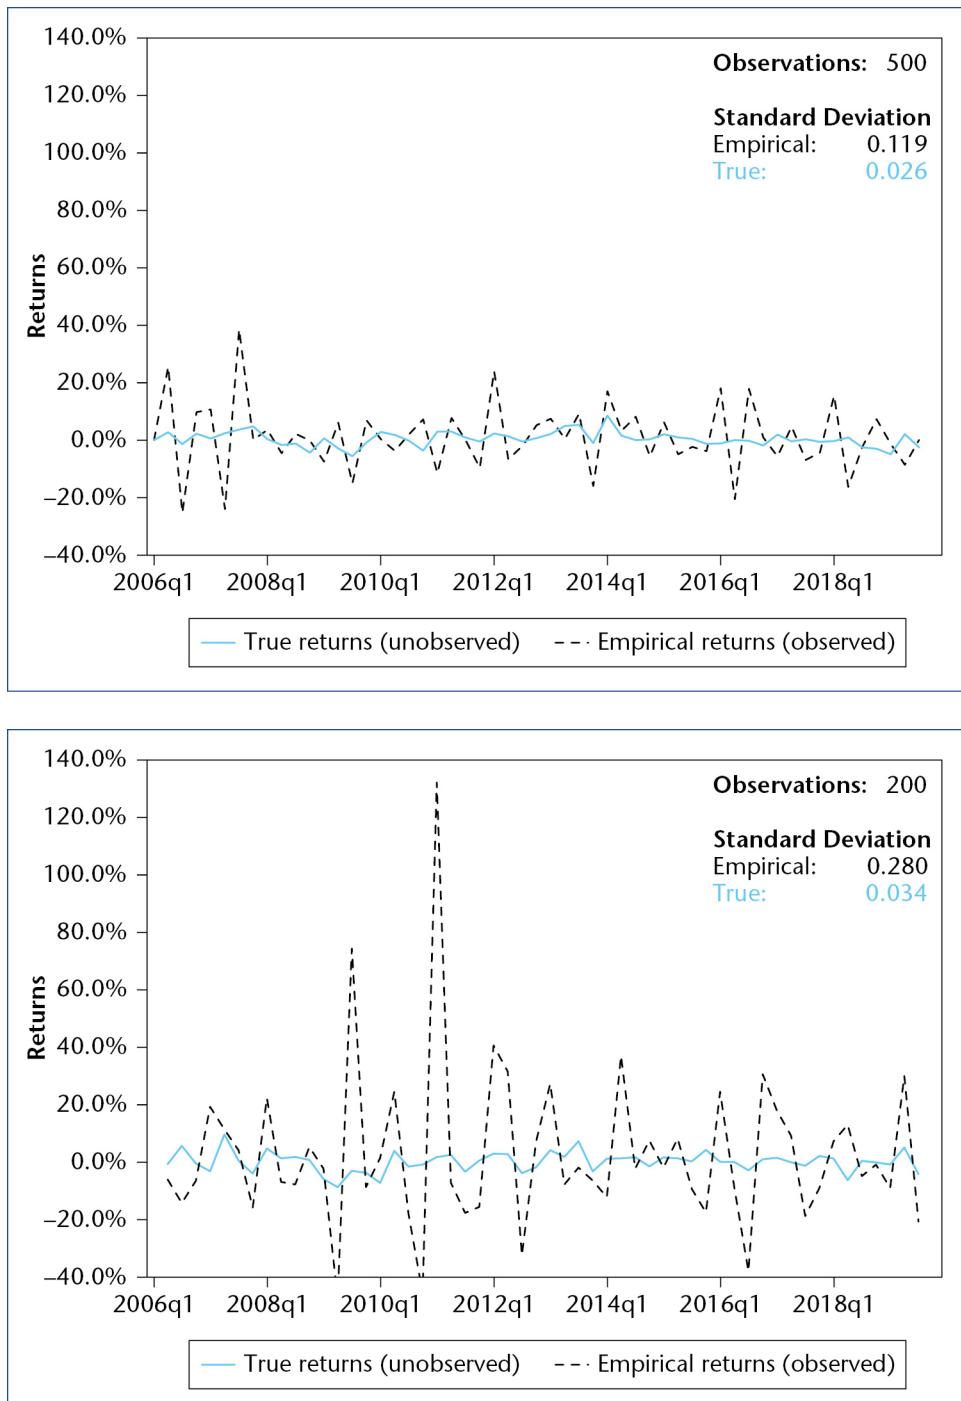


EXHIBIT 24-4A The Pure Effect of Noise in Periodic Returns.

Note: The top is based on a market with 500 observations, the bottom on 200 observations.

Source: Own calculations using data from REZITRADE.

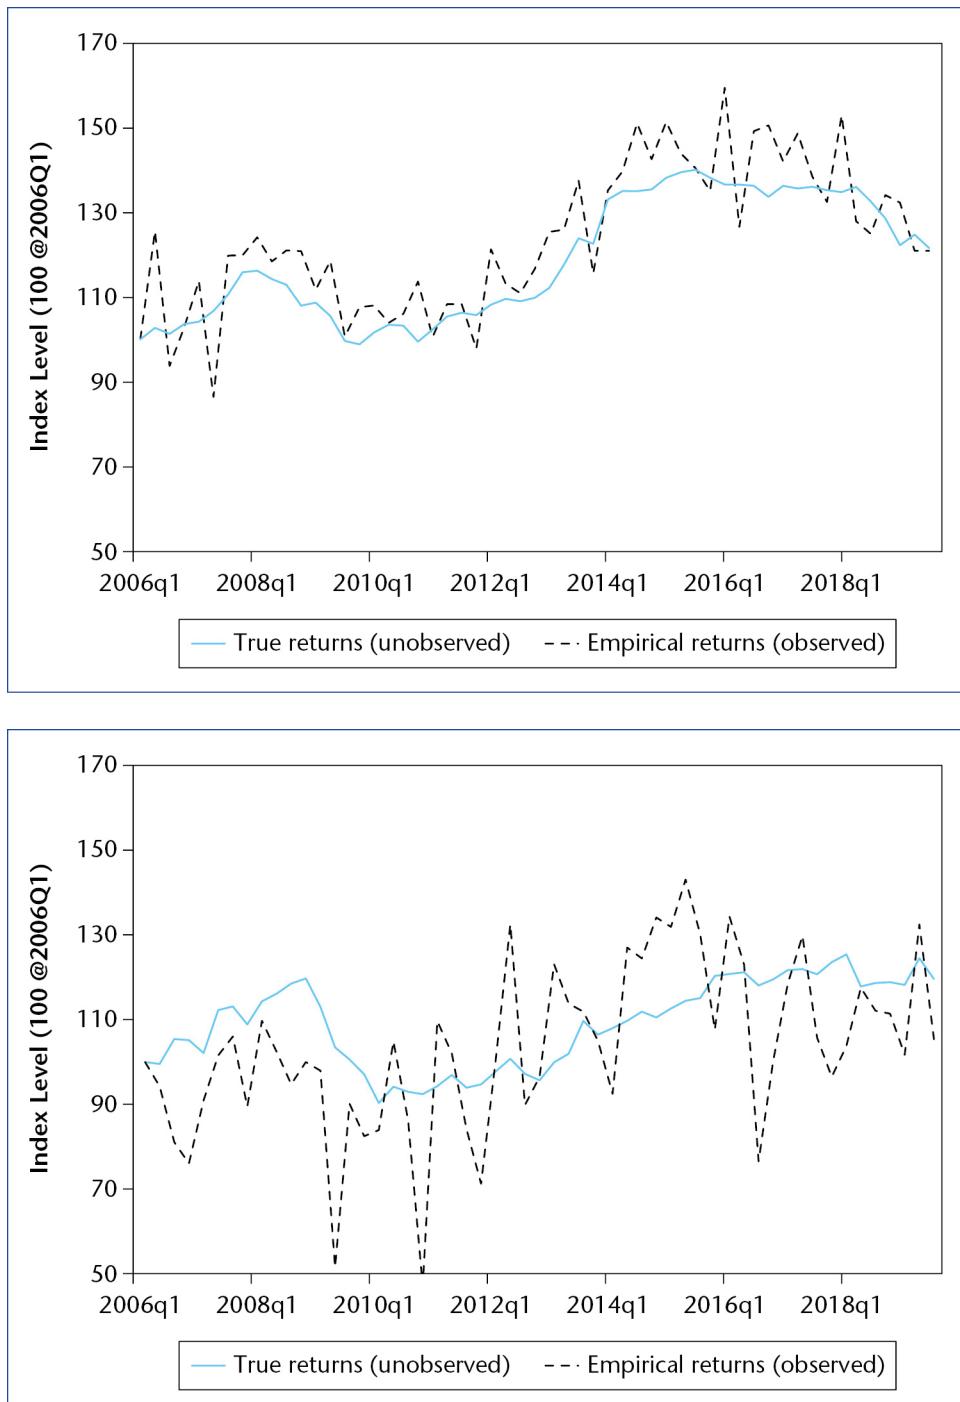


EXHIBIT 24-4B The Pure Effect of Noise in Periodic Value Levels.

Note: The top is based on a market with 500 observations, and the bottom on 200 observations.

Source: Own calculations using data from REZITRADE.

the apparent cross-correlation between the noisy series and any other series.⁹ Thus, noise can make it appear as if two real estate market segments are less correlated than they actually are.

Unlike true volatility, noise does not accumulate over time in the price value levels. Thus, the longer the interval over which returns are measured, the lower the effect of noise on any of the return statistics.

24.2.2 THE PURE EFFECT OF TEMPORAL LAG

Now let's look at the impact on periodic returns of the other type of error common in real estate valuations, temporal lag bias, which does not tend to diversify out at the macro level. To illustrate the pure effect of this type of error, we will consider another unrealistic hypothetical situation. Again, we have a static portfolio, but now we will assume that there are infinite properties in the portfolio so that the effect of random noise completely washes away. However, each year's portfolio valuation is based on appraisals rather than transaction prices. While the appraisals all occur at the end of the calendar year, they display a temporal lag bias in which the appraised value for each property as of the end of year t consists of a one-half weight on the true value as of that time, plus a one-half weight on the true value as of the end of the previous year. The result is that the empirically observed value of the portfolio as of the end of year t , labeled V_t^* is given by the following formula where C_t is the (unobservable) true market value (and we will assume all values are in logs, so that returns are simple arithmetic differences across time):¹⁰

Equation 1:

$$V_t^* = \frac{1}{2}C_t + \frac{1}{2}C_{t-1}$$

The empirically observable return in the index, labeled r_t^* is thus

$$\begin{aligned} r_t^* &= V_t^* - V_{t-1}^* = \left[\left(\frac{1}{2} \right) C_t + \left(\frac{1}{2} \right) C_{t-1} \right] - \left[\left(\frac{1}{2} \right) C_{t-1} + \left(\frac{1}{2} \right) C_{t-2} \right] \\ &= \left(\frac{1}{2} \right) (C_t - C_{t-1}) + \left(\frac{1}{2} \right) (C_{t-1} - C_{t-2}) = \left(\frac{1}{2} \right) r_t - \left(\frac{1}{2} \right) r_{t-1} \end{aligned}$$

Thus, the empirically observable return consists of one-half weight on the true current return plus one-half weight on the previous period's true return. The empirical return is called a **moving average** of the true returns across time.

The effect of return lagging like this is obvious, and it is depicted visually in Exhibit 24-5A and 24-5B, which may be compared to their counterparts in Exhibit 24-4A and 24-4B (left graph, based on the market with 500 observations). The empirical returns in Exhibit 24-5A appear "smoothed" and lagged in time behind the unobservable true returns. The same is true of the value levels depicted in Exhibit 24-5B. Note in particular that the observable value level series tends to be shifted to the right in the graph, with turning points tending to occur later in time, as compared to the unobservable true value level series, and the amplitude of any cycle or rise and fall is slightly diminished.

⁹ The cross-correlation coefficient between two series equals the covariance divided by the product of the volatilities of the two series.

¹⁰ Just to reiterate, there are no random noise terms in equation (1) only because we have "assumed them away," by the device of our hypothetical assumption of an infinite property population in the portfolio. This unrealistic assumption is made to illustrate the pure effect of temporal lag.

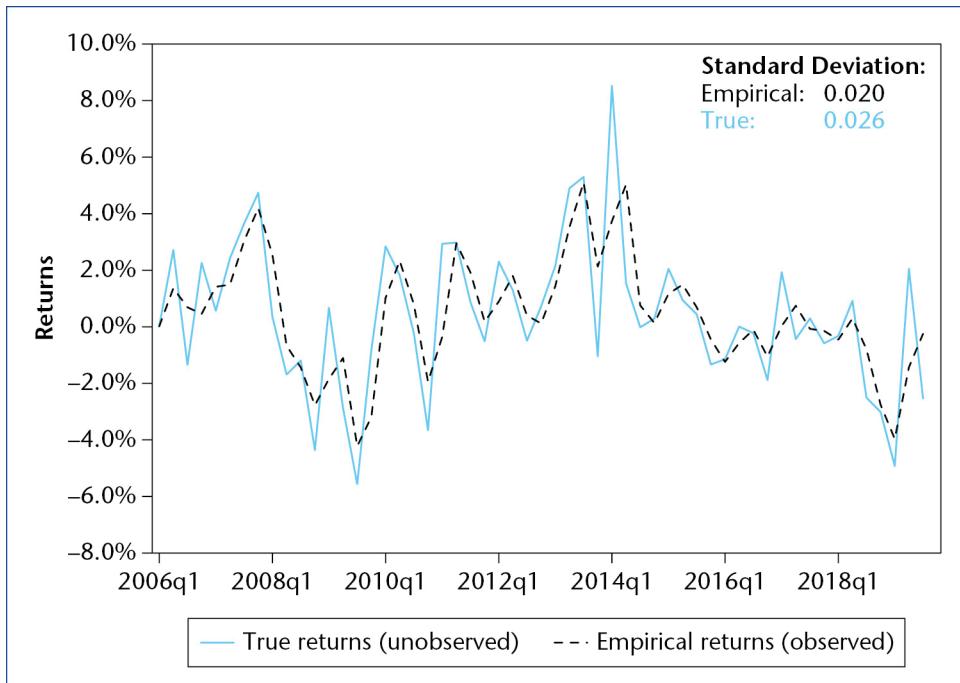


EXHIBIT 24-5A The Pure Effect of Temporal Lag in Periodic Returns.

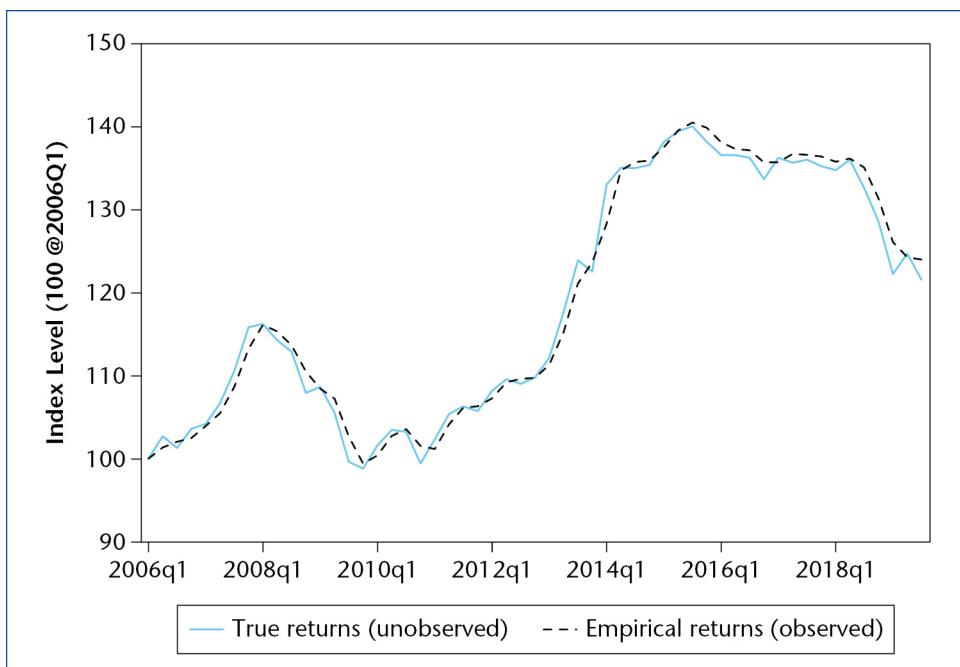


EXHIBIT 24-5B The Pure Effect of Temporal Lag in Periodic Value Levels.

In terms of the effect of temporal lag bias on the periodic returns time-series statistics of interest to investors, first consider the average return across time. Temporal lagging does not change the long-run (or “unconditional”) expected value of the periodic return. However, in any finite sample of time, temporally lagged returns will be “conditionally” biased. That is, the direction of the bias depends on the direction in which the true returns have been trending, if any. For example, if true returns have been increasing over the relevant history, then temporally lagged returns will be biased low, with the average empirical return tending to be lower than the average true return during the history.

Turning to the second-moment statistics, the pure effect of lagging in the form of a moving average of the true returns is to reduce the apparent volatility (total risk) of the observable real estate returns. The more important second-moment effect of lagging, however, is that it reduces the apparent beta (systematic risk) of the real estate returns measured with respect to any nonlagged series, including a nonlagged risk benchmark such as stock market returns. For this reason, the effect of temporal lagging on periodic returns series is often referred to as **smoothing**, or **appraisal smoothing**, as such lagging in appraisal-based indices is often attributed at least in part to the macro-level valuation impact of the micro-level appraiser behavior described in Section 24.1.

What about the cross-correlation and beta of a lagged real estate series with respect to another similarly lagged series? These effects are complicated, considering that the underlying true returns are likely autocorrelated and contain lagged cross-correlation terms. It is likely, however, that the pure effect of temporal lag bias will, in most cases, cause only very slight bias, if any, in the cross-correlation and beta statistics between two similarly lagged real estate series. Thus, for example, the beta of a component of the NPI with respect to the total NPI will probably not be seriously biased, as both series have similar lags.

Regarding cross-temporal statistics, the pure effect of temporal lag bias in the returns series is usually to impart apparent positive autocorrelation into the empirical return series, more so than is present in the unobservable true returns. This tends to exaggerate the effect that long-interval (lower-frequency) periodic returns become relatively more volatile, compared to non-lagged series, than short-interval (high-frequency) returns.¹¹ This causes the apparent long-interval betas of lagged series with respect to non-lagged series to be greater than short-interval betas, and this effect tends to be magnified in the apparent beta of a moving average lagged series. Not surprisingly, moving average lagged returns are more predictable in advance than true underlying returns.

24.2.3 PUTTING THE TWO EFFECTS TOGETHER

In the real world, most practical empirical real estate periodic returns series will contain at least a little bit of both of the “pure” types of errors described in the two preceding sections. The fundamental trade-off between these two sources of error noted in Section 24.1 makes it impossible to completely avoid both error sources simultaneously.

The noise component, n_t , will be more important in smaller portfolios and may well dominate in individual property returns.¹² Because of the Square Root of n Rule, the noise will be less important

¹¹ Both the variance and the covariance with nonlagged series are more than proportional to the length of the return periods in the case of positively autocorrelated periodic returns. For example, annual returns have more than four times the covariance and variance (more than twice the volatility) of quarterly returns, and this effect is magnified the greater is the positive autocorrelation.

¹² Periodic returns series of individual property returns are rarely used as such. In the first place, individual properties are typically rather small assets that have relatively little economic or statistical significance by themselves. Furthermore, in addition to the noise and lag effects described earlier, observable individual property value levels are flat (zero appreciation) between appraisals or transaction observations. This gives individual property returns a very spiky appearance and their value level series a very artificial-looking step-like or sticky appearance. For these reasons (among others), the major return measure used at the individual property level is the multiperiod IRR, rather than periodic HPRs. (See Chapters 9 and 26 for additional discussion relevant to this point.)

in returns to large portfolios or indices. Even though statistical regression is a very efficient tool for handling large amounts of data, noise will still be present in regression-based price indices. Referred to in statistical terms as estimation error, random noise may often be the more important of the two error components in regression-based indices derived from transaction prices, at least provided that the regression is specified and estimated to avoid most temporal lag bias.

In contrast, the temporal lag effect will typically dominate in large portfolios or indices based on appraised values. Temporal lag effects can also be important in regression-based indices if the transaction observations span across time and care is not taken to control temporal aggregation properly in the regression specification and estimation process.

It is important to note that the two types of error effects will tend to mask each other in the empirical returns. In reality, it is impossible to separate the pure effects as we have done in the preceding two sections. For example, the volatility-magnifying effect of return noise will tend to offset, to some degree at least, the volatility-dampening impact of the moving average temporal lag. Similarly, the negative autocorrelation component imparted by the random noise will tend to offset the positive autocorrelation effect of the temporal lag. This type of masking can make it difficult to correct the effect of error on the volatility of the empirical return series. On the other hand, systematic risk (or beta) is more amenable to correction because noise does not affect the theoretical beta; while we know that temporal bias dampens the observable beta toward zero, the more so, the greater the lag.

Exhibit 24-6 presents a visual example of the mixed-errors situation based on the same underlying simulated true returns as those in the previous exhibits. The thin solid line indicates the true returns and market value levels. The dotted line reflects the effect of random noise only. In the real world, we might observe the dotted line in Exhibit 24-6B as the average transaction price in each period among the properties in the subject population that happened to transact during that period.¹³ The dotted line in the returns chart in Exhibit 24-6A is simply the percentage difference in these average prices each period, identical to the returns shown in Exhibit 24-4A.

Now suppose that all the properties in the subject population are appraised at the end of each period, based on appraisers' observations of the transaction prices in the current and past periods (or, equivalently, based on appraisers' partially updating each year, their previous year's appraisal based on the new transaction information). In particular, suppose that appraisers apply the following first-order autoregressive (simple exponential smoothing) valuation equation:

Equation 2

$$V_t^* = \omega \bar{V}_t + (1 - \omega) V_{t-1}^*$$

In particular, suppose that the partial adjustment factor, ω , is 0.2.

Under this assumed appraiser behavior, the thick green line in Exhibit 24-6 traces the appraisal-based index appreciation returns and value levels for our simulated population of properties. Note that the appraisal-based returns and values include both random noise and temporal lag effects, and to some extent, these effects mask each other.

Unlike random noise, the lag of the empirically observable appraised values behind the market values does not tend to get washed out, and it is clearly evidenced in Exhibit 24-6B by the horizontal gap between the green and blue lines in the chart. On average, this gap is about four periods long in our example simulation. It can be seen most prominently in the lagged turning points in the market value cycle and when the observed returns are the most volatile. Of course, this lag would not be directly observable in the real world because the true market returns and values would be

¹³ This would certainly be a finite, and probably even rather small, sample. Therefore, it could not eliminate the noise as we assumed in the preceding section when we assumed an infinite population.

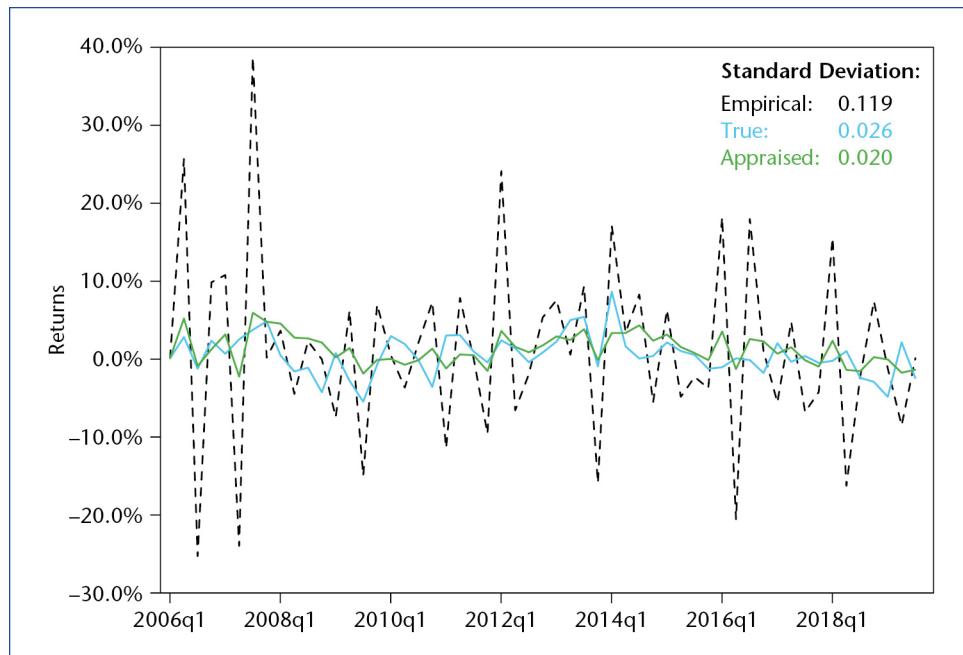


EXHIBIT 24-6A Periodic Appreciation Returns Based on Market Values, Transaction Prices, and Appraised Values.

Notes: The appraised returns (V^*) are calculated as $V_t^* = 0.2\bar{V}_t + (1-0.2)V_{t-1}^*$, where \bar{V} are empirical returns (i.e., coming from transaction data). See Equation (3) as well.

unobservable. The same lag also exists between the appraisal-based index and the average transaction prices, which would be empirically observable. However, the noisiness of the transaction price series obscures the picture, making it difficult to see clearly the lag in the appraised values relative to the transaction prices.

24.2.4 WHAT ABOUT THE TOTAL RETURN?

The preceding three sections introduced you to the major types of measurement errors that are unique to private real estate regarding empirical periodic returns series. The discussion and examples so far considered only the appreciation component of the total HPR. However, most of the volatility in periodic returns series derives from the appreciation return component, so the points made previously about appreciation returns generally carry through to the total returns as well.¹⁴ Furthermore, accurately observing the income returns is often much more straightforward (at least for properties in an investment database such as NCREIF, IPD, or for REITs). So, the primary challenge in tracking private property market returns is measuring the asset price or capital return component.

¹⁴ Remember that the denominator in the formula for the periodic income return component is the beginning period asset valuation. Thus, income return components can be affected directly by valuation errors, particularly in appraisal-based indices. For example, assuming contemporaneous income is accurately observed, the income return component will be biased low when the asset valuation is biased high, and vice versa. These effects will generally be minor, however, at least relatively speaking.

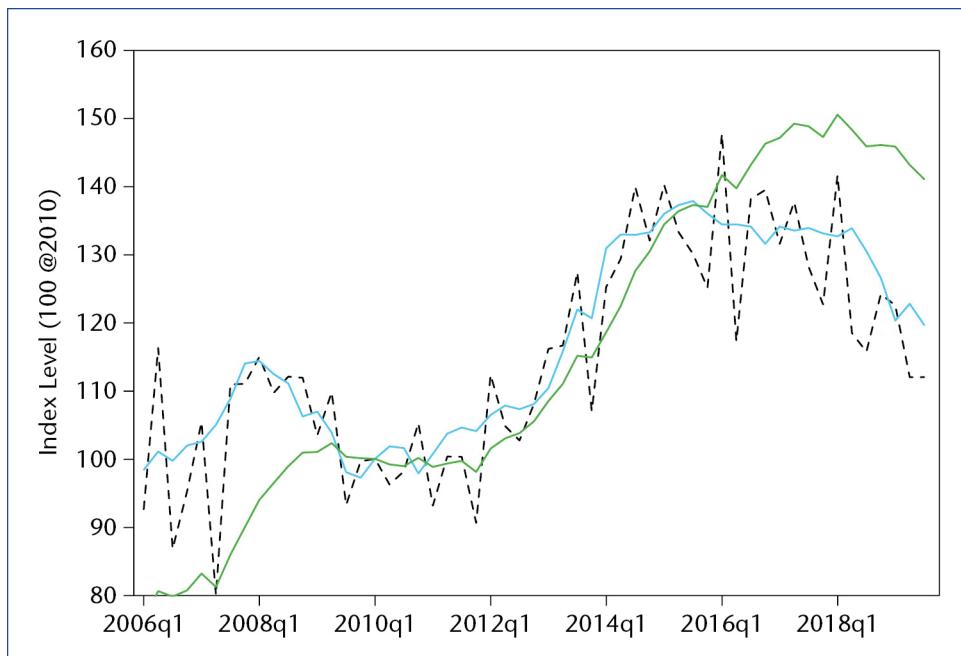


EXHIBIT 24-6B Market Values, Transaction Prices, and Appraised Values.

Notes: The appraised returns (V^*) are calculated as $V_t^* = 0.2\bar{V} + (1 - 0.2)V_{t-1}^*$, where \bar{V} are empirical returns (i.e., coming from transaction data). See Equation (3) as well.

24.3 WHAT IS TRUTH? LAGS AND THE TIMELINE OF PRICE DISCOVERY IN REAL ESTATE

The previous section made clear that temporal lag bias is one of the major types of errors in private real estate periodic returns data. This type of error can be particularly prominent and important in macro-level appraisal-based returns, such as those tracked for large portfolios and benchmark indices such as the NPI. With this in mind, we need to step back and consider a broader and deeper question.

What do we mean by the “true return”? After all, it is impossible to define *error* without defining a specific conception of the correct value. As a result, it is impossible to appreciate the nature and magnitude of the temporal lag bias problem in real estate without such a definition.

Up to now, we have been defining true returns as those based on market value, as this term has been previously defined in this book regarding the private real estate asset market, namely, the *expected* (or *ex-ante*) sale price in the current relevant market for the subject asset. However, the problem of temporal lag bias often arises in contexts that are different from or broader than what is implied by this definition of value. For example, consider the following practical decision problems, all of which involve, directly or indirectly, the use of appraisal-based periodic returns data for private real estate:

1. A private market real estate investment manager using a timing-based tactical investment policy wants to know if the institutional office property market “turned around” last quarter, or was it the previous quarter, or two quarters ago, or has it not yet turned?
2. A private market real estate investment manager wants to hedge some of its real estate market risk exposure by taking a short position in a real estate index swap derivative.

3. A consultant considering portfolio allocation strategies wants to know the long-run beta (or long-run relative volatility and correlation coefficient) between real estate and the stock market.
4. An advisor to a wealthy individual trying to decide between private direct investment and REIT shares for the real estate component of her wealth portfolio wants to know the beta of private real estate compared to that of REITs in both up and down stock markets.
5. An acquisition officer involved in a protracted negotiation for a major property wants to know how far property market values have fallen overall during the past calendar quarter.
6. A manager of a large portfolio whose annual incentive fee is pegged to the NPI is wondering why his recently appraised portfolio didn't beat the NPI last year, even though he believes his appraisers are competent. He's sure he did better than most of his competitors last year.
7. An appraiser wants to know the approximate ex-post time-weighted mean total return risk premium of institutional quality real estate over T-bills during the past quarter century to help her estimate the appropriate cost of capital for a DCF valuation.
8. A pension fund wants to obtain risk and return exposure identical to the NPI by taking a long position in an index swap contract based on the NPI.

For the managers in problems 1 and 2, a definition of market value based on contemporaneous empirically observable transaction prices in the private market, such as would be tracked by a transactions-based index, will be best. The consultant in problem three might find something closer to a REIT market-based definition of value more relevant for defining or measuring the temporal lag that is (or should be) of most concern in her problem of measuring long-run beta, at least from the perspective of investors who want to use the public market as their real estate investment vehicle. The advisor and the acquisition officer in problems 4 and 5 might prefer a constant-liquidity definition (see next Section 24.3.1) of private market value rather than one based only on consummated transaction price evidence. The relevant value for the manager in problem six might be a fully contemporaneous appraisal-based valuation of his benchmark. Finally, for the appraiser and fund manager in problems 7 and 8, an appraisal-based index that is not entirely contemporaneous, like the NCREIF index, would be quite adequate (or indeed, for the fund manager in problem 8, the NPI specifically is the stated preference).

If you do not have a clear conception of the relevant value of the problem you are trying to address, you can get very confused in your attempts to analyze the problem. Yet this question of "truth" (or, what is the relevant conception of value) is so basic and underlying in nature that analysts and decision-makers often do not even explicitly consider it. The result is practical decision situations in which the analysts are confused without even knowing they are confused, a potentially dangerous situation! To avoid this situation, we think it is helpful to consider a timeline of **price discovery** as it relates to different types of empirical and conceptual real estate periodic returns series.

24.3.1 MULTISTEP REAL ESTATE TIMELINE OF INFORMATION INCORPORATION INTO OBSERVABLE VALUE

Consider the following simple situation to build your intuition regarding different relevant "true values" in real estate. Suppose that for a long time, no new information has arrived relevant to the value of real estate assets. Prices of REIT shares, as well as prices of privately traded property, not to mention appraised values, are all stable and steady. The time is now $t-1$, and all real estate price indices are at a level we can set at 100.

Now, into this very dull world, a piece of news, that is, new information, arrives suddenly, unpredicted (of course, or it wouldn't be news). The information arrives precisely at time t and is relevant to the value of all real estate assets. In fact, for illustrative purposes, suppose that this information is

relevant to all real estate assets in the same direction and magnitude (although this will not become fully apparent empirically for some time). In particular, as a result of this new information, the value of all real estate assets has just decreased by 10 percent. Then, immediately after this one piece of news arrives, the world becomes very boring again, and no new information relevant to real estate values arrives for a long time.

OK, it's a strange and very unrealistic situation we have just depicted, but it serves to illustrate the timeline we want you to become familiar with. Let's consider how different real estate prices or value indices might respond to the news that arrived at time t . In particular, we will identify and define five different such indices at a conceptual level: (1) An index based on REIT share prices, (2) a constant-liquidity private market value index, (3) a contemporaneous transaction-price-based private market value index, (4) a contemporaneous appraised value index, and (5) an appraisal-based index with staggered appraisals or some "stale" appraisals such as the NPI. We will consider each of these in turn as we walk through the index response pattern illustrated in Exhibit 24-7.

Index 1: REIT Share Price-Based Index. This index is based on the market prices of publicly traded REIT shares on the stock exchange. Most directly, this would include REIT indices such as the NAREIT Index in the United States. It could also include property-level indices (of REIT property assets de-levered from the REITs' financial structure) targeted at specific property sectors, such as the FTSE NAREIT Real Estate Index Series. The stock market where REITs trade is the densest, most liquid market relevant to trading real estate equity assets. As a result, it has the most informationally efficient price discovery. The REIT-based index moves first and fastest. It always reflects always liquid prices, in the sense that investors can always sell at the prevailing price. However, Index 1 may be subject to some excess volatility, based perhaps on "herd behavior" or overshooting by investors in the REIT shares in the stock market. This is illustrated in Exhibit 24-7 as we depict the REIT index overshooting the mark at index level 90, falling slightly below that level before it corrects itself.¹⁵ Line 1 in Exhibit 24-7 traces the path of the REIT index over time from just before the arrival of the news to past the time when all indices have fully reflected the news. In practice, REIT indices show relatively little autocorrelation or inertia. However, there may be some minor positive autocorrelation in very short-run returns, followed by a tendency for slight negative autocorrelation at longer frequencies (presumably reflecting correction in the excess volatility noted previously).

Index 2: Constant-Liquidity Private Market Value Index. This is an index of market values in the private property market. However, this index simultaneously reflects both trading volume and transaction prices. Volume is as important an indicator as consummated prices for tracking a whole-asset search market like real estate. The constant liquidity index represents the asset values that would equilibrate the market with a constant amount of liquidity over time. When bad news arrives, news that suggests market values should fall, real estate market participants become cautious. Potential buyers reduce their reservation prices while property owners on the supply side either don't reduce their reservation prices as much or may even increase them (reflecting uncertainty about the market conditions), causing the two sides of the market to pull

¹⁵ The existence of this type of excess volatility in the stock market is somewhat controversial, although some tendency in this regard seems to be fairly widely accepted now among financial economists, and there are several theories to explain it (e.g., irrational investor behavior models, noise trader models, and so forth). Excess volatility of this nature has much in common with the random noise described in Section 24.1.1 regarding private real estate indices, as far as some statistical effects are concerned. However, stock market excess volatility is different in its basic nature and source from the type of noise we considered previously, as it does not generally imply the existence of any measurement error in the market values or returns that are quantified in the index. Furthermore, in some circumstances, private property market values bring relevant information to the REIT market, as indicated by the widespread use of NAV data in analyzing REIT stock values. Such information can help correct the overshooting that REIT share prices are prone to in the stock market.

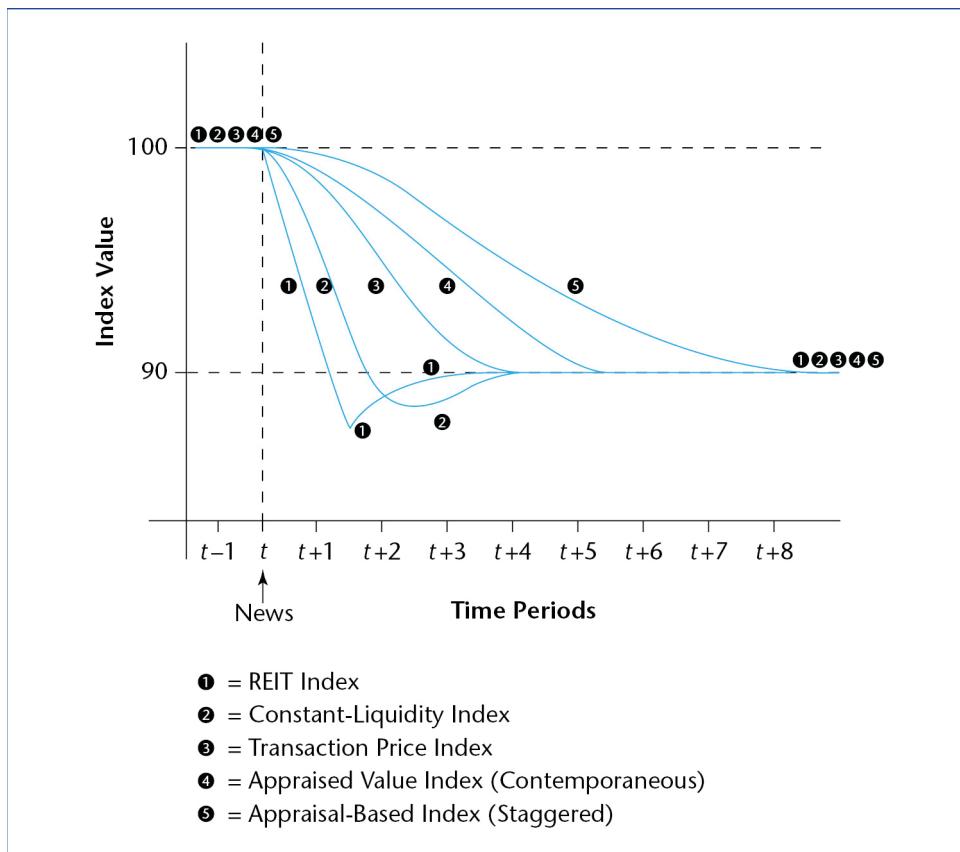


EXHIBIT 24-7 News-Response Timeline for Various Stylized Real Estate Indices.

away from each other, reducing liquidity and transaction volume in the market.¹⁶ In contrast, when the market turns up, buyers tend to move farther or faster than sellers, bringing reservation prices closer together (with more overlap - recall Exhibit 24-1) and increasing liquidity and transaction volume in the market. Thus, liquidity (as reflected by trading volume turnover) in the private market is “pro-cyclical,” varying positively with the market cycle. In order to maintain constant liquidity, in other words, to maintain a constant expected transaction turnover ratio, seller reservation prices would have to track those of buyers. An index tracking the buyer (demand) side movements in the property market would therefore tend to reflect greater volatility or faster response to news than an index based on consummated transaction prices. In response to our hypothetical single-news event, the typical price path of such a constant-liquidity private market index is indicated by line 2 in Exhibit 24-7. Note that it lags behind the stock market-based Index 1 and, like Index 1, might tend to overshoot slightly the ultimate value change.

¹⁶ Recall that reservation prices are the prices at which potential participants in the market will stop searching or negotiating and commit to a deal. This same type of uncertainty happens in the REIT market. Still, it gets resolved much more quickly due to the much denser market, with multiple active buyers and sellers publicly trading homogeneous shares with very low transaction costs and the ability to engage in short-selling. Increased uncertainty may be associated with increased volatility in the REIT market.

Index 3: Contemporaneous Transaction-Price-Based Index of Private Market Values.

This is an index reflecting the central tendency (or average) of the sales prices of the actually consummated transactions within each period of time. This type of index can be estimated directly using empirically observable transaction prices. The two major methods to compute such indices are the hedonic and the repeat-sales regression-based models. The latter is more widely used, as repeat-sales indices reflect the actual round-trip price experiences of investors in the market (see Section 24.4.2). Examples of this index in the United States include MSCI's Commercial Property Price Indices (CPPI) and the CoStar Commercial Repeat-Sales Indices (CCRSI).¹⁷ Of course, any empirically based transaction price index would, in practice, tend to exhibit some estimation error or transaction price noise of the type described in Section 24.2.1. However, recent developments in index methodology have minimized the impact. Since Index 3 represents prices that can reflect widely varying liquidity (transaction volume) over time in the private real estate market, it must be kept in mind that the returns reported in this index do not reflect a constant ease of selling or a fully complete picture of the state of the market in the absence of corresponding trading volume information. Liquidity variation is one reason we would expect consummated transaction prices in the private property market to move slightly behind or smoother than the constant liquidity values tracked by Index 2. Thus, at least at the conceptual level (in the absence of noise), transaction price indices like Index 3 should be a bit less volatile and slightly more lagged in time as compared to the two previously defined indices, as suggested by line 3 in the Exhibit. However, in an important sense, Index 3 is the most fundamental indication of property value. The prices it reflects represent the prevailing equilibrium in the property market, the actual prices paid and received by investors on both sides of the market (supply and demand).

Index 4: Contemporaneous Appraised Value Index. This is an index based on micro-level appraised values, that is, the cross-sectional aggregation of individual property appraisals, but with all of the properties being appraised each period as of the current point in time. Thus, conceptually, the only difference between this index and the transaction-price-based index is the temporal lag present in optimal micro-level appraisal, as described in Section 24.1 (as in the difference between points B and C in Exhibit 24-3). This index will thus be a bit smoother and lagged in time behind the transaction price index as appraisers wait to compile evidence from transaction prices before they finalize their own opinions of how the news that arrived at the time changed the value of real estate. Ultimately, an appraisal-based index will settle for the same illustrative "true" value of 90 as the previously described indices, but only a little later. Exhibit 24-7 indicates the contemporaneous appraisal-based index by line 4.

Index 5: Appraisal-Based Index with Staggered Appraisals. This index is based on the same type of valuation observations as those in the previous index, only in this index, not all properties are reappraised in each period that the index is reported or as of the same point in time within each period. This results in the valuations reflected by the index being partially "stale" (or out of date) and/or spread out over time (which tends to smooth and lag the aggregate value in the index). This makes Index 5 a lagged moving average of the contemporaneous appraisal-based index, as depicted in line 5 in Exhibit 24-7.

Let us now summarize the timeline we have just described, as illustrated in Exhibit 24-6. The effect of the news arriving at time t is incorporated at varying rates into the real estate prices or values defined and measured in the five different ways described by these five conceptual indices. The index value levels traced through time from the left-hand edge of Exhibit 24-7 at the old value level of 100 are splayed out and arrive at the new ultimate value level of 90 spread across time. The

¹⁷ We do not cover the major private data vendors here aside from mentioning CoStar and MSCI, but others that supply transaction data include CREXI, RealPage, ICE (Black Knight), RealNex, Reonomy, and many others.

REIT, constant liquidity, and contemporaneous transaction price index arrive at the new valuation first.¹⁸ The contemporaneous appraised values follow them. Then the appraisal-based index with the staggered (stale) appraisals pulls up the rear, arriving last at the 90 mark.

Obviously, the temporal lag bias in real estate values depends on which of these five indices one is using, and which one is taken as the truth against which the bias is defined and measured. Each of the five indices represents a different way of measuring real estate value and, therefore, in a sense, a different *definition* of value. Suppose we use a staggered appraisal-based index such as Index 5 (e.g., an index similar to the NCREIF Property Index - NPI). In that case, the bias is relatively small if a contemporaneous appraisal-based index such as Index 4 defines the relevant true value. However, the temporal lag is much longer, and the smoothing is much greater when measured or compared against a REIT-based index such as Index 1 or a constant-liquidity transactions-based index such as Index 2.

What do these different conceptual indices of real estate value movements look like in empirical reality? This answer will depend on specific data and estimation methodologies employed to construct the indices and will no doubt differ across countries, reflecting different market and appraisal practices and cultures. Exhibit 24-8 shows the historical price levels for four regularly published indices as of 2012 representing the institutional investment commercial property market in the United States during the 2000-2011 period of history (which, of course, included a great bull market followed by a major financial crisis and then a recovery).¹⁹

Exhibit 24-8 shows only four of the five conceptual indices defined in Exhibit 24-7, omitting Index 4 (the contemporaneous appraisal-based index).²⁰ The representative for Index 1, the stock market-based index, is the FTSE NAREIT Real Estate Index. Being based on REIT share prices, this index publishes daily updates. It is published in both equity and property level versions, with the latter being de-levered so as to track the stock market's valuations of the REITs' property assets directly. The other three indices depicted in Exhibit 24-8 are all based on the NCREIF property database, representing pension funds and other nonprofit institutions' investments in commercial property. The authors have estimated the constant liquidity index in Exhibit 24-8 (corresponding to Index 2 in Exhibit 24-7) using NCREIF data.²¹ The transactions-based index reflects the contemporaneous transaction prices of the properties sold from the NCREIF Index each quarter.²² Finally, the NCREIF Property Index (NPI) is based on the official valuations reported to NCREIF each quarter by the data-contributing investment manager members of NCREIF. The NPI is the oldest (going

¹⁸ Probably in that order, that is, the REIT index first, then the constant-liquidity index, then the transaction index.

¹⁹ During COVID, the response was not as clean as it was around the GFC. Mostly because the response was very different depending on the type of property during COVID. For example, initially, values for suburban apartments and industrial warehouse space went up, while office space saw a big price drop. (Retail prices were already going down before COVID.) In contrast, everything went down equally during the GFC. Thus, the waters are a bit muddied in making pre/post-COVID comparisons across indices.

²⁰ There is not a good representative of Index 4 in the United States, though the NPI has in recent years approached more closely to that archetype as the frequency of appraisals has increased, especially among the ODCE Index properties (the ODCE is a NCREIF fund-level index of open-end, diversified core equity funds whose properties are appraised frequently).

²¹ The exact estimation is out of the scope of this book. For further reading see Fisher, J., D. Geltner, and H. Pollakowski, 2007, "A Quarterly Transactions-based Index of Institutional Real Estate Investment Performance and Movements in Supply and Demand," *Journal of Real Estate Finance and Economics*, 34: 5–33.

²² The NTBI was first developed in 2006 at the MIT Center for Real Estate with NCREIF cooperation. The index was initially based on a hedonic type of regression model using the recent appraised values of the sold properties as the major explanatory variable for the transaction prices, with time-dummy variables to capture differences each quarter between the transaction prices and the appraisals (to correct for the lagging and smoothing bias). In 2011 the TBI was taken over by NCREIF and the methodology was slightly changed, so that the NTBI became a SPAR (sales price / appraisal ratio) based index in which the average ratio of sales price to recent appraised value across the transactions in each period was multiplied by the appraisal-based index value each period to derive the implied transaction price levels. The NTBI was discontinued in 2020 as the COVID pandemic greatly reduced the number of property transactions relative to the number of properties in the NPI.

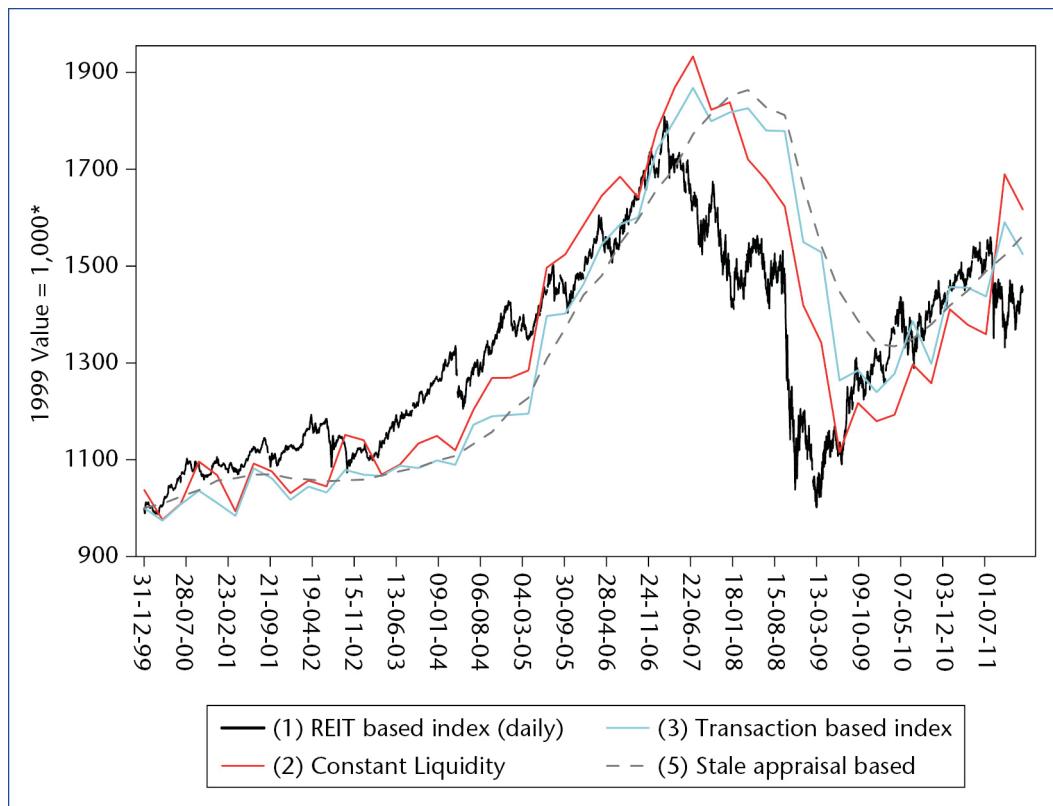


EXHIBIT 24-8 Four Definitions and Measures of U.S. Commercial Property Values, 2000–2011.

Note: REIT-based index is from NAREIT (real estate index), whereas the other indices are based on NCREIF data. NCREIF publishes index (5) under the name NPI.

* The only exception is the constant liquidity index, which is demeaned around the transaction-based index.

Source: Author's compilations from NAREIT and NCREIF.

back to 1977) and most traditional index of U.S. institutional investment property performance. In recent years, NCREIF members have begun reappraising their properties more frequently and with less seasonal bunching in the fourth quarter, which has enabled the NPI to be more up-to-date than in previous decades.

The indices in Exhibit 24-8 are based on real-world empirical data, and so include some noise and much more complex behavior than the stylized pure-type indices depicted in Exhibit 24-7 in a simple world where only one piece of information arrived. However, the general patterns and relationships depicted in Exhibit 24-7 nevertheless reveal themselves in Exhibit 24-8. All four indices trace a similar history at the broad-brush level, including the major boom, bust, and recovery. The stock market-based Index 1 moves first and crashes a bit farther, only to experience a greater or faster recovery. It also seems to show greater short-run or transient volatility, although part of this may simply be that it is updated so much more frequently (daily instead of quarterly for the other three indices). In fact, measured on quarterly-frequency returns, the FTSE-NAREIT Real Estate Index quarterly volatility during the 2000-2011 period was about the same as that of the transaction-based index (Index 3).

The solid red line representing the constant-liquidity Index 2 moved ahead, and more exaggeratedly, compared to the direct transaction price Index 3 (dotted blue line). This difference reflects

(and/or causes) the extreme liquidity cycle that occurred in the private property market during the period shown. When the financial crisis struck in 2008 and early 2009, potential investors on the demand side of the market drastically dropped their reservation prices (the constant liquidity index lost 42 percent of its peak value from the end of the second quarter of 2007 through the second quarter of 2009). However, property owners held back and did not reduce their reservation prices nearly so much or so fast. The result was a drastic drop in trading volume. Transactions among NCREIF properties declined by more than 60 percent, from over 10 percent of all properties in 2007 to only 4 percent in 2009.

Measured on a calendar quarter basis, the magnitude of the property price crash from each index's respective peak to trough (whose dates can be seen visually in the Exhibit) varied from 41 percent in the REIT Index to 42 percent in the constant liquidity index, to 34 percent in the transaction-based index, to 28 percent in the appraisal-based NPI. Through the end of 2011, the subsequent recoveries were (as a percent of the trough value), respectively: 36 percent, 45 percent, 26 percent, and 17 percent.

24.4 INNOVATIONS IN COMMERCIAL PROPERTY RETURNS INDICES

In this section we will first summarize the major recent developments in indexing commercial investment property prices and returns in the United States. Then we will discuss in more depth a particularly important innovation, that of repeat-sales indices of transaction prices.

24.4.1 OVERVIEW OF TYPES OF INDICES

From the previous sections in this chapter, you can see that there are potentially three major different types or sources of data for constructing indices to track commercial property prices and investment returns over time:

- Independent professional appraisals of individual properties
- Direct transaction price evidence from the property market
- Stock market-based evidence from REIT share prices.

Indices based on each of these three types of data have different characteristics, strengths, and weaknesses.

Appraisal-based indices have been the only sort of index available until recently in the United States and still in most other countries. They can be very useful, as they are typically based on a well-defined population of properties whose values and income are reported into the index database using standardized, professionally agreed-upon criteria. So, appraisal-based indices tend to be very well documented. They can support detailed performance attribution as well as drill-down into submarkets, and their databases generally allow measurement of total investment return (including net cash flow), not just price change or capital return.²³

On the other hand, appraisal-based indices may suffer from the smoothing and lagging bias noted earlier in this chapter, and they are relatively subjective or “self-reported” by the property owners or managers, a consideration that can sometimes be an issue. Furthermore, in the United States the appraisal-based NCREIF and IPD indices represent a rather narrow and specialized population of properties (though an important one): those owned by pension funds and other investment

²³ Performance attribution was introduced at the micro-level in the supplemental materials of Chapter 4 and will be discussed at the macro-level in Chapter 25. Also, we noted earlier in the present chapter that, while the income component of the total return is quite important, it is not the main challenge or major focus of commercial property investment performance indexing. The main focus here is on tracking asset price changes over time.

institutions that must “mark to market” their property values frequently.²⁴ The broader population of commercial investment properties is not regularly appraised and cannot be tracked by an appraisal-based index.

As discussed in Chapter 1, the total value of all investable commercial real estate is around \$18 trillion. These are characterized by being owned or financed by large investment institutions such as major banks, insurance companies, pension funds, private equity funds, REITs, CMBS, and foreign investors and are traded relatively frequently. Within this population, the NPI only covers some 16,000 properties worth approximately \$1.1 trillion, a small fraction of the total.

REITs own \$2.5 trillion in properties (as of 2024). Since REITs are effectively “pure plays” (essentially owning nothing but commercial property investment assets), they can indirectly provide another important source of information about commercial property values via their stock market share prices. More specifically, the enterprise value of a REIT (stock value plus debt) can be seen as the stock market’s valuation of the underlying real estate. However, as we describe in Chapter 26, the stock and private property markets do not always “agree” on the value of commercial property. Currently, most commercial properties in the United States can only be tracked by direct transaction price evidence from the private property market. Another downside of using REIT data is that although REITs tend to specialize by property type, they do not do so by geographical location. Thus, it is difficult/impossible to obtain local price dynamics from REIT data, as REITs own properties in multiple locations.²⁵

Transaction price evidence is arguably the most fundamental type of evidence about property value, as noted in our previous discussion of Index 3 in Exhibit 24-7. Appraisers depend heavily on observations of transaction prices in order to make their valuation judgments. As property assets actually trade in the private property market, their transaction prices in that market most fundamentally underpin and reflect the actual experiences of real estate investors. With this in mind, let us now focus on the major innovation in commercial property indices of the past decade: the advent of repeat-sales-transaction-based indices of commercial property price performance.

24.4.2 REPEAT-SALES INDICES OF COMMERCIAL PROPERTY PRICE PERFORMANCE

Until recently, there were insufficient databases of commercial property transaction prices in the United States to compile a price index based on them. But by the turn of the present century, CoStar, Real Capital Analytics, and a few other companies had developed large, sophisticated, and constantly updated databases of commercial transactions. This, combined with advances in econometric and computational technology for producing price indices in private whole-asset markets, enabled the production and regular publication of statistically rigorous indices of commercial property prices.

Given a good database of asset prices, it is easy to tabulate the average or median price per square foot each period. Such indices have long been produced for the housing market in the United States. But commercial property is much more heterogeneous than housing, and there will generally be a serious “apples-versus-oranges” problem if you take a simple comparison of the average price in one period, compare it to the average price in the preceding period, and try to infer from that comparison a price change or capital return as experienced by investors. For example, some average price indices increased during the market downturn from 2008 to 2009. But that was only reflecting a “flight to quality,” as nervous investors shifted their activity toward buying nothing but the highest

²⁴ As U.S. GAAP accounting standards, which generally use historical cost, evolve more towards international standards that rely more on fair value, it may be that broader populations of properties in the United States will begin to be frequently and regularly appraised. This may allow an expansion of the population of properties that could be tracked by appraisal-based indices.

²⁵ Recent literature has tried to fill this void, see among others, Geltner, D., Kumar, A., & Van de Minne, A. M. (2023). “Estimating Commercial Property Fundamentals from REIT data,” *The Journal of Real Estate Finance and Economics*, 1–23.

quality, most prime properties. Even in the housing market, the leading price indices now used in the most sophisticated analyses are based on statistically rigorous procedures that formally control for the apples-to-oranges problem of different houses trading in different periods. The most widely used such procedure is the repeat-sales price index.

With the **repeat-sales model**, you hold constant *all* the characteristics of the properties being sold in every period using regression techniques (Chapter 9). As a result, the change in prices over time is purely due to changes in the market conditions, *not* changes in the composition of properties sold. This is arguably the most relevant measure of interest to investors.²⁶ See the text box for how to build your own repeat sales index.

TEXT BOX 24.1 BUILD YOUR OWN REPEAT SALES MODEL

In Chapter 9, we discuss the inner workings of regression models. It allows you to hold certain characteristics constant and estimate elasticities between variables. However, as noted in the printed text, a repeat sales model controls for *all* characteristics. How does this work? First, we would have to discuss the term **dummy variable**. (Which we acknowledge is a funny word.) A dummy variable can only take up a value of 1 or 0. It is 1 when the observation falls in a specific category and zero otherwise. Such a category could be “*sold in 2020*.” As you can see, this is either true (1) or not (0). Suppose you have a **transaction data set** that reports (i) the location or unique ID of the property, (ii) the year a property was sold, and (iii) the sales price. In the U.S., deeds report all three and are publicly available in all but a few states. Say we find seven transactions in the last 3 years for three unique properties (Main St, Public Pl, and Ajax Sq) in our relevant market. Consider the following data format.

Transaction	Main Street #1	Public Plaza #2	Ajax Square #3	Sold in 2022	Sold in 2023	Sold in 2024	Price/Sqf
1	1	0	0	1	0	0	\$ 200
2	1	0	0	0	1	0	\$ 240
3	0	1	0	1	0	0	\$ 100
4	0	1	0	0	1	0	\$ 100
5	0	1	0	0	0	1	\$ 110
6	0	0	1	1	0	0	\$ 400
7	0	0	1	0	0	1	\$ 500

As you can see, a property is either located on Main Street #1 or not. The first two transactions are for that address. You can think of columns 2 through 4 as “address dummies.” The next three columns provide the year in which the property was transacted. The property at Main Street #1 was sold in 2022 and 2023 (but not 2024). These are called the “periodic dummies.” The final column provides the price per square foot. Next, you regress the price

²⁶ There may be other foci of interest. For example, instead of focusing on investors’ experiences, one could focus on the average price trend in the market. Though closely related to investors’ experiences, the market average price trend is not exactly the same thing as the average investor’s price change experience. Investors buy and sell at specific times at least partly of their own choosing, and they must always sell the same property that they have previously bought. In contrast, the market exists continuously, and its stock of buildings is generally being always updated and rejuvenated with new buildings, such that the same buildings do not characterize the market across time. For example, the average age of the building stock in the market does not increase year for year like the age of any given building. Other types of price models and indices may be more appropriate than repeat-sales for measuring the market price trend, at least in principle, such as hedonic price modeling. We put a primer on statistics for real estate online with this book. The primer will show you how to do your own hedonic model with a small data sample.

per square foot on the “address dummies” and “periodic dummies.” The online Excel file will provide all the necessary steps in more detail. The address dummies will essentially demean the prices per square foot *per property*. Hence, the periodic dummies will capture the “pure” change in price per square foot from one period to the next, as separate from the change in composition. If done correctly, you will find that prices rose by almost \$60 per square foot during this time period.

However, an increase of \$60 per square foot in two years is not yet an index. A price index is ultimately used to produce price returns. A \$60 increase is a lot for the \$100 per square foot building (60 percent increase), but a lot less for the \$500 per square foot building (12 percent increase). It is, therefore, common to log transform sales prices. The reason is that the periodic dummies can be interpreted directly as a percent price change *common for all properties*. An extra benefit is that one does not need to take the price per square foot anymore. As long as the property did not change its square footage between transactions, you can take the log of the price directly. In many cases, the square footage is unknown or reported inconsistently.

Unfortunately, estimating standard repeat sales models (like in the textbox) has resulted in noisy “indices that are too “noisy.” This is especially the case for smaller markets with fewer observations. Hence, providers of standard repeat sales models in the early 2000s stuck to the national level. However, having more granular indices is *more useful* for many obvious reasons. Indeed, the price dynamics and risk/return profiles can vary considerably between markets and property types. However, as you go more granular, the indices become noisier and hence *less useful*. We showed this effect earlier in Exhibit 24-4. In fact, we did not note it then, but the “empirical” series was based on a repeat sales model using actual data. Note that even at as many as 500 observations, we have a relatively noisy index that does not pass the “sniff test.” (Or rather, “eyeball test.”) For commercial real estate, it is not uncommon to have less than 100 transactions in smaller markets. Also, in certain quarters, there might not even be a transaction at all. The standard repeat sales model cannot produce a return in such a situation.

Unsurprisingly, in the 2010s, considerable advancements have been made to extract the true signal from such noisy real estate time series. This form of signal extraction, which falls under what is called *structural time series models*, is quite literal rocket science. It was developed first as part of the Manhattan Project in the 40s and was later used by NASA to get the first people on the moon. Even though the math behind such models is complex, the intuition is rather straightforward. For example, say we would like to know how much prices went up for Boston office in the last month. We have two sources we can rely on: (i) transaction data and (ii) a panel of experts. Suppose we have two observations in our transaction data for the last month. The first transaction indicates that prices went up by 50 percent, and the second transaction shows that prices went down by 30 percent, for an average of -10 percent price appreciation, according to the data. Next, the panel of experts shows up. When asked for their prediction, *all* of them responded that prices went up by 5 percent. What is the truth in this case?

Under structural time-series modeling, the answer would be close to +5 percent. The reason is that there was *no disagreement* between the panel of experts, and the transaction data was all over the place. The reverse could also happen. If we have many data points, and they all point in the same direction, such models would put more weight on the data. Essentially, what happens is that we take a weighted average of the prediction and the data based on the (bell) shape of the distribution of both. A wide distribution (high standard deviation) means more uncertainty and, thus, less weight. The natural follow-up question is, where do I find this large group of experts? We admit we bamboozled you here, as these do not exist. In fact, we ask a model to provide a prediction. Such models are typically similar to the VAR introduced in Chapter 9! Note from Chapter 9 that the VAR provides both a mean prediction and an uncertainty around this prediction. This mean and

uncertainty are overlaid with the mean prediction and uncertainty (i.e., the distribution) that comes from the data to provide a weighted mean between the two.

Structural time-series models allow for highly accurate and less noisy indices. To return to our example of Exhibit 24-4, the blue line, which we called the “true” index, was actually based on such a model. As shown in those Exhibits, that index looks legitimate even with just 200 observations. Obviously, in reality, we cannot know with certainty that these are the “true” indices. However, we know that such models are closer to the “truth” than standard repeat sales models. (Both academics *and* industry agree on this.)

With all of this in mind, let’s look at the recent price-change histories evidenced in the first regularly published structural time series repeat-sales indices of commercial property prices in the United States. The first was the MSCI Commercial Property Price Index (CPPI), developed originally by Real Capital Analytics in 2016. This index has an inception date as of the end of 2000 and tracks properties generally over \$2,500,000 in value. This threshold roughly distinguishes the “professional investment” segment of the commercial property market from the smaller properties that typically do not trade as frequently and are often occupied and used by their owners. By now, MSCI publishes 100s of CPPIs across property types and markets for a variety of countries.

Let’s start with some higher aggregate-level office indices in the United States. Exhibit 24-9 provides the National, “Major Market,” and “Non-Major Market” quarterly indices for the office market. The Major Markets are New York, Los Angeles, Washington, D.C., Chicago, San Francisco, and Boston. Thus, this index is estimated based on office transactions in all those markets. The

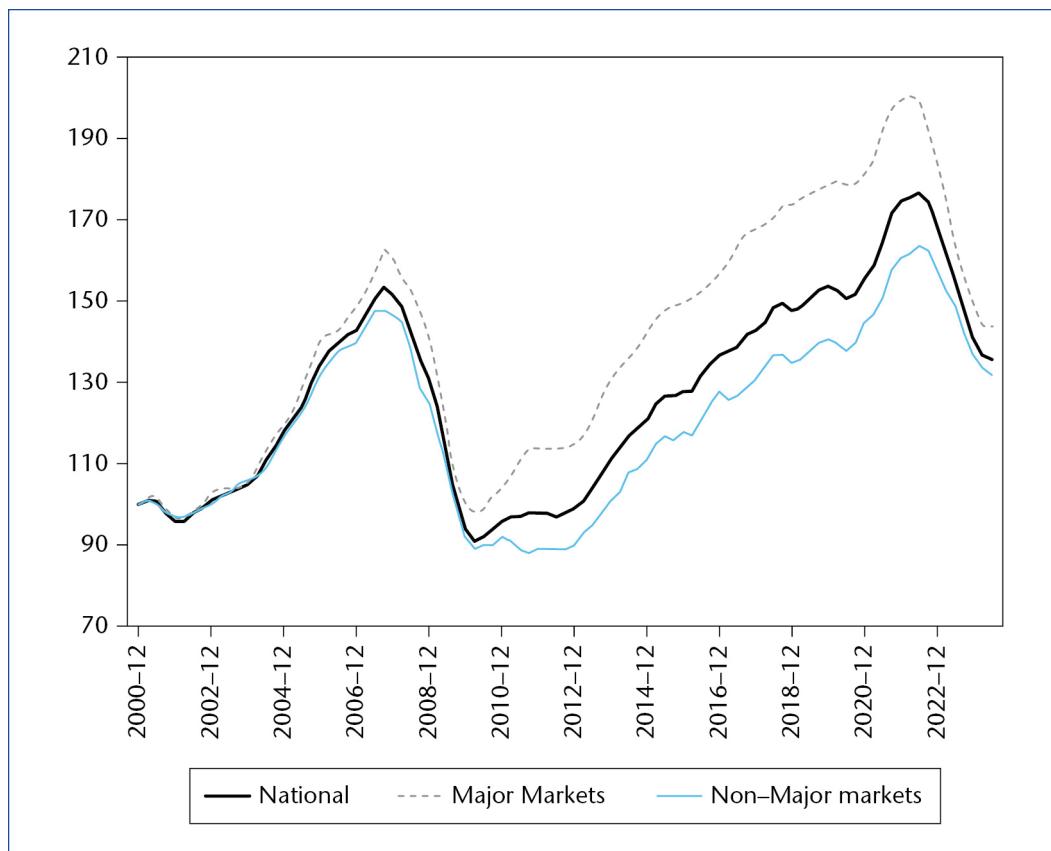


EXHIBIT 24-9 National Office indices.

Source: MSCI CPPIs.

Non-Major Markets consist of office transactions elsewhere. Keeping in mind that these indices represent only the price-change relevant to the investors' capital returns, not the total investment returns, it is clear that offices in Major Markets had seen higher price increases compared to their Non-Major Markets counterparts in times when things went well. For example, just before COVID-19 hit, the office index for the Major Markets sat at 200, whereas it was at 160 for the Non-Major Markets index. At the same time, the Major Markets index also falls more during times of distress. For example, after COVID, the Major Market office index dropped by 25%, whereas the Non-Major office index dropped by 18 percent. Something similar happened during the GFC.

Although there are already some differences at the national level, the larger differences are at a more granular level. In Exhibit 24-10, we provide a sample of such indices. For the more granular markets, MSCI clusters office, industrial, and retail into one group called "commercial," or COM for short. (This is different from Apartments or APT.)

In the Exhibit, the random assortment of indices shows some similarities. All indices were clearly impacted by the GFC (albeit the timing was different). There are also many differences. For example, the three apartment indices were the clear winners at the end of the sample (2024). However, COVID-19 has had a larger impact on apartments compared to commercial properties. This doesn't tell the full story, though. First of all, by combining office, retail, and industrial into one group, the impact of COVID partly cancels out. We know that office properties were impacted

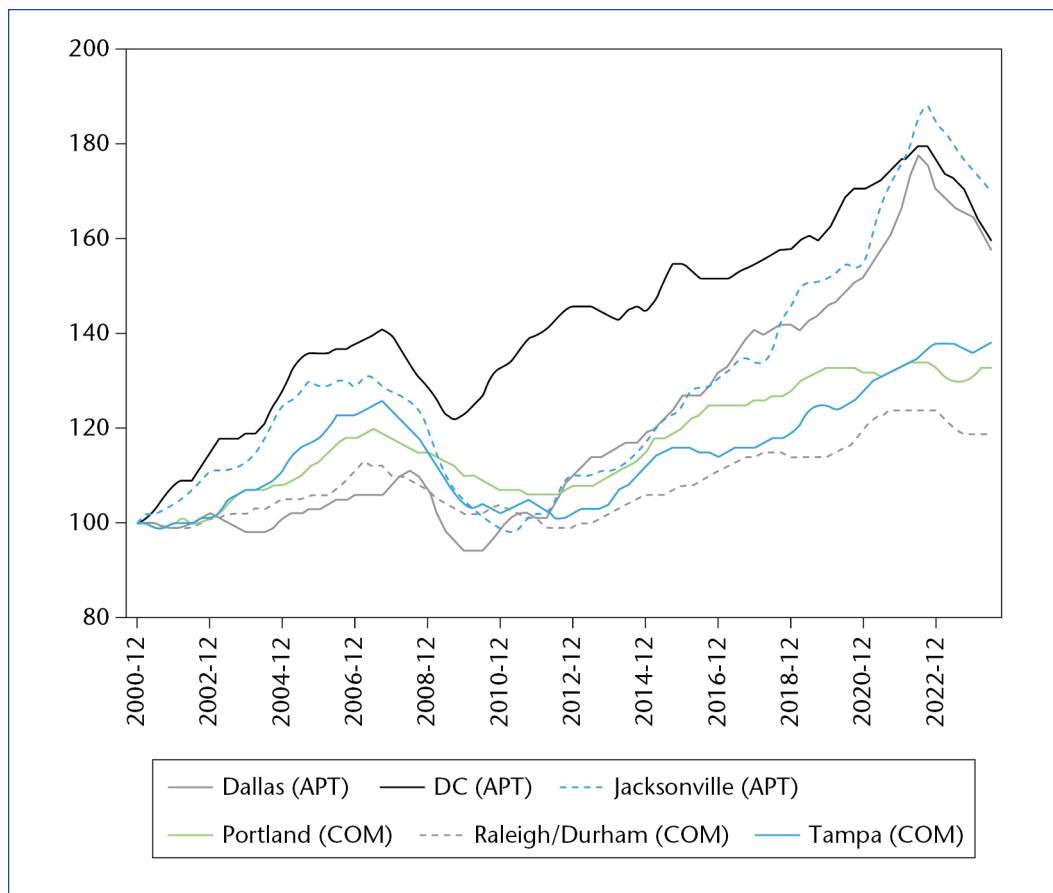


EXHIBIT 24-10 Sample of Indices for Commercial Real Estate.

Source: MSCI CPPIs.

negatively; however, industrial properties were impacted positively (as there was more demand for goods that were shipped directly to the home). Secondly, liquidity indices showed declines that were magnitudes larger compared to transaction price changes during COVID, especially for offices. This indicates another “flight to quality,” as was the case during the GFC. We added a short academic paper to the online supplemental materials examining such liquidity indices during COVID-19.

In short, despite the industry taking great steps, obvious improvements still need to be made. For example, the unbundling of office, industrial, and retail for more granular markets and the need for constant liquidity indices, especially during times of distress. As of this writing, companies like MSCI, Moody’s, and even the Federal Reserve are working on exactly those types of products. The field of property price and return indexing is still rapidly developing, and the next decade may see further important innovations.

24.5 CHAPTER SUMMARY

In this chapter, we have provided an in-depth discussion of the nature of private market real estate, periodic price change, and investment returns data—the raw material on which virtually all quantitative macro-level real estate investment analysis is based. In Chapter 25, we will address some subjects that practitioners in the real world of macro-level real estate investment deal with on a daily basis. You will see again how the data issues described in this chapter are relevant.

KEY TERMS

- NCREIF Property Index (NPI)
- Error
- Random noise
- Square Root of n Rule
- Temporal lag
- Volatility effects (of return errors)
- Covariance effects
- Moving average
- Smoothing (or appraisal smoothing)
- Beta effects
- Cross-correlation effects
- Autocorrelation effects
- Price discovery
- Index 1 (REIT-based returns)
- Index 2 (constant-liquidity private market returns)
- Index 3 (observable contemporaneous transaction price returns)
- Index 4 (contemporaneous appraisal-based returns)
- Index 5 (staggered appraisal-based index returns) noise-versus-lag trade-off repeat-sales price index

STUDY QUESTIONS

Conceptual Questions

- 24.1. What are the two most common or prominent types of “error” in major indices of private real estate price appreciation?
- 24.2. Considering appraisal valuations of individual properties, explain the meaning of random error (or noise) component of appraisal “error.”

24.3. In the same context as above, what is meant by the temporal lag bias component of appraisal error? What is the typical cause or source of such bias? What is the implication of temporal lag bias for the distribution of appraised values around the underlying true value?

24.4. How can it be rational for an appraiser to only partially adjust her opinion of value of a property in response to the arrival of new information?

24.5. What is the noise-versus-lag trade-off in individual property appraisal? How is this related to the Square Root of n Rule? Which of these two types of error is relatively more important at the micro-level of individual property value estimation? What about at the macro-level of static portfolio value estimation as the aggregation of the appraised values of individual properties in the portfolio?

24.6. Describe and contrast the pure effect of random noise and the pure effect of temporal lagging in terms of how an empirical periodic price index would differ from the unobservable underlying true returns. When is the noise effect relatively more important? When is the temporal lag effect typically more important?

24.7. What risks arise when examining the role of real estate in mixed-asset portfolios without correcting for biases in appraisal-based real estate return series, especially when determining optimal asset class allocations that include real estate?

24.8. In what way can random noise and temporal lag bias mask each other in typical real world price indices? When is such masking likely to be most important?

24.9. Section 24.3.1 presents five different conceptual definitions of commercial property value levels (or periodic appreciation returns).

- Describe each of the five value indices, starting first with Index 5, which is a staggered appraisal-based index like the NCREIF property index (NPI), and then moving up to Index 1, which is stock market based index such as the FTSE-NAREIT PureProperty Index. As part of your answer, explain the advantages and disadvantages of each index relative to the previous one you just described. What do you gain and what do you lose?
- Describe the dynamic behavior of each value definition in response to the arrival of news relevant to the value of real estate assets. (See Exhibit 24-7.)
- Which of the indices do you think is the most relevant for each of the eight practical decision problems posed at the beginning of Section 24.3? Explain your choice in each case.

24.10. Explain why the dynamic adjustment of a transaction-based (variable liquidity) real estate price index is more sluggish than a constant liquidity version of the index. How is this related to the fact that prices and trading volume are jointly determined in the private real estate asset market? Based on your answers, explain why constant liquidity prices were higher than transaction-prices in the mid-2000s boom years, but lower during the downturn in 2008-09, as seen in Exhibit 24-8.

24.11. What are the three major types of price or investment performance indices publicly available for tracking private commercial property in the U.S. as of the early 2010s?

24.12. What is the essential nature of a repeat-sales price index? That is, in what way does it essentially reflect the experience of investors in the private property market?

Quantitative Problems

24.13. Three nearly identical properties are all sold as of the same date, for prices of \$2,650,000, \$2,450,000, and \$2,400,000.

- What is your best estimate of the market values of each of the two other properties that are virtually identical to the first three, as of that same date?

- b. If one of those additional properties actually did sell for a price of \$2,550,000, what is your estimate of the market value of the fifth property as of that same date?
- c. What is the market value of the fourth property?
- d. What is your best estimate of the error or noise (defined as the difference between the transaction price and the market value) in the observed transaction prices of all four properties?

24.14. In the previous problem, how much did the accuracy of your estimate of market value improve as a result of the addition of the fourth empirical valuation observation (i.e., what is your estimation error with the fourth transaction, and what is it with only the first three)?

24.15. Suppose you regress a time-series of appraisal-based index periodic returns onto both contemporaneous and lagged securities market returns that do not suffer from lagging or measurement errors. That is, you perform the following regression, where $r_{M,t}$ is the accurate market return in period t and r_t^* is the appraisal-based real estate return in period t :

$$r_t^* = \alpha + \beta_0 r_{M,t} + \beta_1 r_{M,t-1} + \beta_2 r_{M,t-2} + \beta_3 r_{M,t-3} + \mu_t$$

The resulting contemporaneous and lagged beta values are

$$\begin{aligned}\hat{\beta}_0 &= 0.05 \\ \hat{\beta}_1 &= 0.15 \\ \hat{\beta}_2 &= 0.10 \\ \hat{\beta}_3 &= 0.00\end{aligned}$$

What is your best estimate of the true long-run beta between real estate and the securities market index?

24.16. Suppose you have reason to believe that appraisal behavior is well characterized by the following autoregression model relating quarterly appraised values to average transaction prices:

$$V_t^* = 0.30\bar{V}_t + 0.70V_{t-1}^*$$

where the V values are in log levels. Now suppose that a quarterly appraisal-based index (based on fully contemporaneous appraisals, that is, without stale appraisals) indicates an appreciation return of 1% in the current quarter (t), and 1.5% in the previous quarter ($t-1$). Assuming that the appraisal-based index has been cleansed of random noise, what is the implied current period contemporaneous transaction- price-based appreciation return (like Index 3) suggested by a reverse engineering of the appraisal behavior?

25 Real Estate Investment Management

CHAPTER OUTLINE

- 25.1 Macro-Level Investment Performance Attribution
 - 25.1.1 The Importance of Property-Level Investment Management
 - 25.1.2 Macro Property-Level Performance Attribution and the Four Fundamental Responsibilities
 - 25.1.3 Portfolio-Level Performance Attribution
 - 25.1.4 Use of a Benchmark in Performance Attribution
 - 25.1.5 The Case for Using Manager Allocation Weights in the Benchmark
- 25.2 Investment Performance Evaluation and Benchmarking
 - 25.2.1 The Basic Idea
 - 25.2.2 Benchmarks and Investment Styles in the Private Real Estate Asset Class
- 25.3 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- What is meant by investment performance attribution at both the macro property level and the portfolio level.
- How to quantify segment allocation versus asset selection effects in a portfolio's differential performance relative to an appropriate benchmark.
- What are the major institutional real estate investment management "styles" and how these may be accounted for in quantitative investment performance evaluation.

Recall that at the outset of Part VI we highlighted three major types of macro-level real estate investment decisions: strategic policy, tactical policy, and policy implementation. In this last chapter of Part VI we turn our primary attention to this last question: how to implement real estate investment policy, or more broadly, how to manage real estate investments. This subject is obviously important for those who will be doing the management directly themselves, but it is also important for those whose money will be managed by others with more specialized expertise in real estate management. Decision makers at both levels should be familiar with the principles and tools we will introduce in this chapter. The chapter will begin with macro-level investment performance attribution and manager evaluation. Finally, we will go into more depth to examine some nuts and bolts of the investment management business with a consideration of benchmarking and real estate investment styles. (More depth and detail are available in the two chapter appendices accessible online in the chapter supplemental materials folders.)

25.1 MACRO-LEVEL INVESTMENT PERFORMANCE ATTRIBUTION

In the field of sophisticated, institutional real estate investment management, there is growing potential to apply quantitative analytical tools to help understand and improve investment performance. A fundamental metric that can be useful in this regard is investment **performance attribution**. The idea is to break down into components the ex-post total return achieved by a given portfolio or manager in order to learn something about the nature and source of the total return performance. Such dissection can then be combined with **benchmarking**, the comparison of the subject manager's or portfolio's performance results with the average results of some suitably defined index or peer group of similar properties or managers, to gain further insight.¹ Performance attribution is useful for two broad purposes: diagnosis and evaluation. The former purpose is primarily of interest to managers, to help them understand the sources and causes of their realized performance. The use of performance attribution in manager evaluation is of interest both to managers and investor clients, as it can help investors to judge and compare the performance results of competing managers.

25.1.1 THE IMPORTANCE OF PROPERTY-LEVEL INVESTMENT MANAGEMENT

Two salient features of investment in the private real estate market as distinguished from most investment in publicly-traded securities (including REITs) are that: (1) The investor cannot simply passively invest in an index but rather must buy (and hold and ultimately sell) specific individual property assets; and (2) the investor is responsible for the operational management of those assets.

With this in mind, look at Exhibit 25-1. The thick line in the exhibit depicts the Moody's/RCA Commercial Property Price Index (CPPI) from its inception at the end of 2000 with a value of 100 through June 2012, when this index tracking same-property prices had a value of approximately 140.² The thin gray lines show the price performance of the 10 properties from this index that were last sold in 2012. Each property is pegged to the index value at the date when it was purchased, and then its value relative to the index is "grown" at a constant rate per month so as to reveal its cumulative price change differential relative to the index by the time of its sale in June 2012. The exact historical data and endpoints in time do not matter. The point is to show graphically how individual property investment performance can differ dramatically from that of the average or aggregate market as a whole, which is represented by the index. For example, one of those 10 properties ended up more than 50 percent below the aggregate index value, while another couple of properties ended up close to 50 percent above the index. From a statistical perspective, this type of property-level performance dispersion is random, both in time and in direction and magnitude. But as an investment manager responsible for buying, operating, and selling specific individual properties, how can you try to make sure your properties are among the "winners," or at least not among the worst "losers"? The answer to this question can be significant in determining how successful you are as a real estate investor in the private property market, even apart from what happens with the market as a whole and the average investor.

25.1.2 MACRO PROPERTY-LEVEL PERFORMANCE ATTRIBUTION AND THE FOUR FUNDAMENTAL RESPONSIBILITIES

(See also our 5-minute Video 1 in the Chapter 25 Supplements folder.)

In Appendix 4A at the end of Chapter 4, we described *micro-level* performance attribution, an approach to "parsing" (or breaking down into its components) the IRR at the level of the individual

¹ Quantitative performance attribution in real estate investment was pioneered to a considerable degree by the British company, Investment Property Databank (IPD), starting in the 1980s. As late as the 2010s such formal analysis was still more prevalent in the United Kingdom than in the United States, though it is becoming more of a standard state of the art practice among many global real estate investment management firms and funds, and both IPD and NCREIF collect and disseminate data relevant for performance attribution and benchmarking in the United States.

² In between, as you can see, the index peaked around 180 in 2007 and then fell to barely 110 in 2009. See Section 24.4.2 in the preceding chapter for a description of repeat-sales indices, the type of price index the CPPI is. (Note that the exact

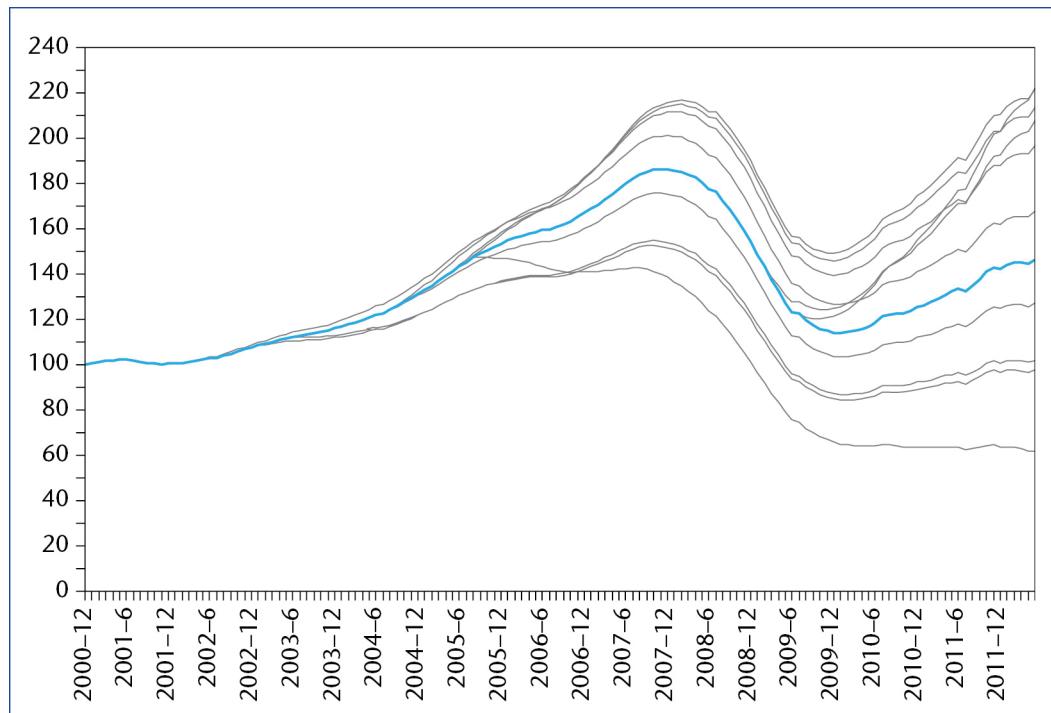


EXHIBIT 25-1 Dispersion of 10 Individual Property Price Changes Around an Aggregate Index (Moody's/RCA CPPI): 10 Last-Sold Property Investments From the Index, Taken as of a Random Point in Time (June 2012).

Note: Pegged to starting value on index at time of prior sale (buy), then tracked at constant rate per month to ultimate sale price deviation from index at disposition in June 2012.

Source: Data from Real Capital Analytics Inc. now MSCI (used with permission).

property investment so as to provide some insight regarding the sources of the overall multiperiod return on the investment. At what might be termed the **macro property level**, performance attribution can also apply the same IRR parsing as was described in Appendix 4A, only aggregated across an entire portfolio of properties. The objective is to attribute the property-level since-acquisition IRRs of the investments to: (1) the initial cash yield, labeled **Y**, (2) the cash flow change during the holding period, labeled **CFC** and (3) the effect of the change in the cash yield between the acquisition and the disposition (or current appraisal if the property is still being held), labeled **YC**.³

Macro property-level performance attribution is most useful to the investment manager for diagnostic purposes, to help the manager understand the property-level sources of his performance results (e.g., not counting the effect of leverage or capital structure⁴), particularly if the manager's

values in the CPPI change as new data is incorporated into the compilation of the average investors' round-trip realized price-change experiences, as described in Chapter 24.)

³ See Appendix 4A in Chapter 4 for further details and a numerical example. There are also other ways of defining and measuring property-level performance attribution, most notably that promulgated by MSCI/IPD. The key feature of the system defined in the Appendix 4A is that it is based on the multiperiod IRR since the acquisition of each property, including the effect of the acquisition price. This sharpens the focus of the measure onto the round-trip investment performance at the property level, including all four major property-level investment management functions of property selection, acquisition, management during the holding period, and disposition (if the property has been sold).

⁴ Performance attribution can in principle be applied to the levered equity performance, or the effect of leverage can be isolated first and treated separately, by comparing the levered equity performance versus the property-level performance without leverage.

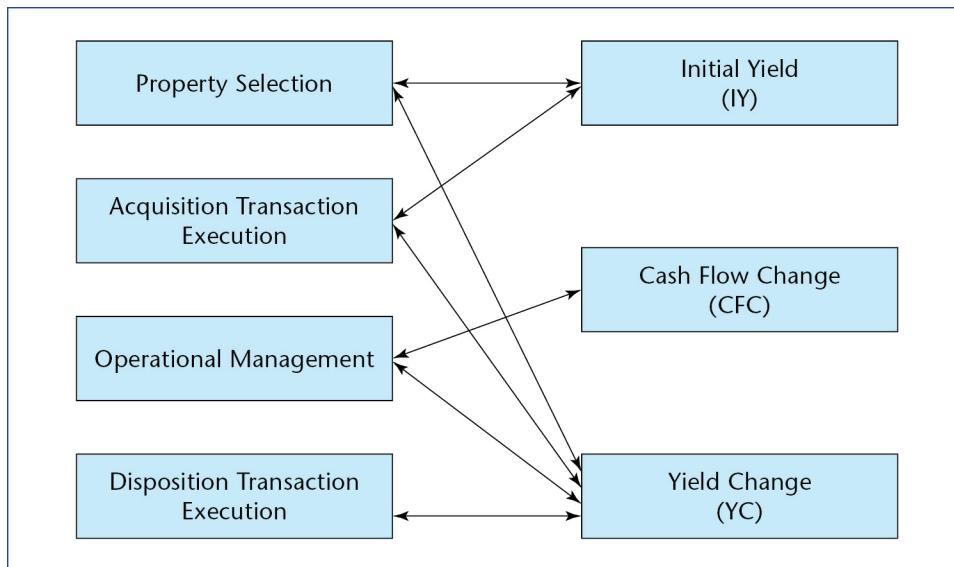


EXHIBIT 25-2 The Four Fundamental Property-Level Investment Management Responsibilities Related to the Three Property-Level Performance Attributes.

results are compared with an appropriate benchmark of other similar properties owned by competing investment managers in the subject manager's peer universe.

The basic idea in property-level since-acquisition IRR performance attribution is to relate the investment performance to the *four fundamental responsibilities* of property-level investment management:

1. **Property selection:** Picking "good" properties (e.g., bargains as found).
2. **Acquisition transaction execution:** Skillful negotiation and due diligence.
3. **Operational management during the holding period:** Including property-level marketing and positioning, leasing, tenant servicing, expense management, profitable capital expenditure management, and so on.
4. **Disposition transaction execution:** Successful timing, effective search for the right buyer, skillful and efficient sales management, and negotiation.

Property-level since-acquisition IRR attribution can help an investment manager understand their strengths and weaknesses at the property-level because the four fundamental responsibilities relate differentially to the three IRR attributes, as suggested in Exhibit 25-2. Thus, how well the manager scores quantitatively in the three attributes reveals something about relative strengths and weaknesses, or at least raises questions which suggest avenues for analysis and perhaps improvement. The links depicted in Exhibit 25-2 are of course not exact or mechanistic. Other factors can intervene to affect performance, such as market movements beyond the control of the manager. Performance attribution is an art, not a science. Nevertheless, careful application of property-level performance attribution analysis may be able to provide insight. It can also help to set an expectation of accountability within the investment management organization.

As is typical with all performance attribution, property-level performance attribution is often most effective when applied in conjunction with benchmarking. The idea is to compare the investment performance of a subject property or portfolio or manager with that of an appropriate cohort of similar properties over the same period of time. For core (stabilized, institutional quality) investment properties in the United States, the NCREIF index can be used to synthesize a benchmark,

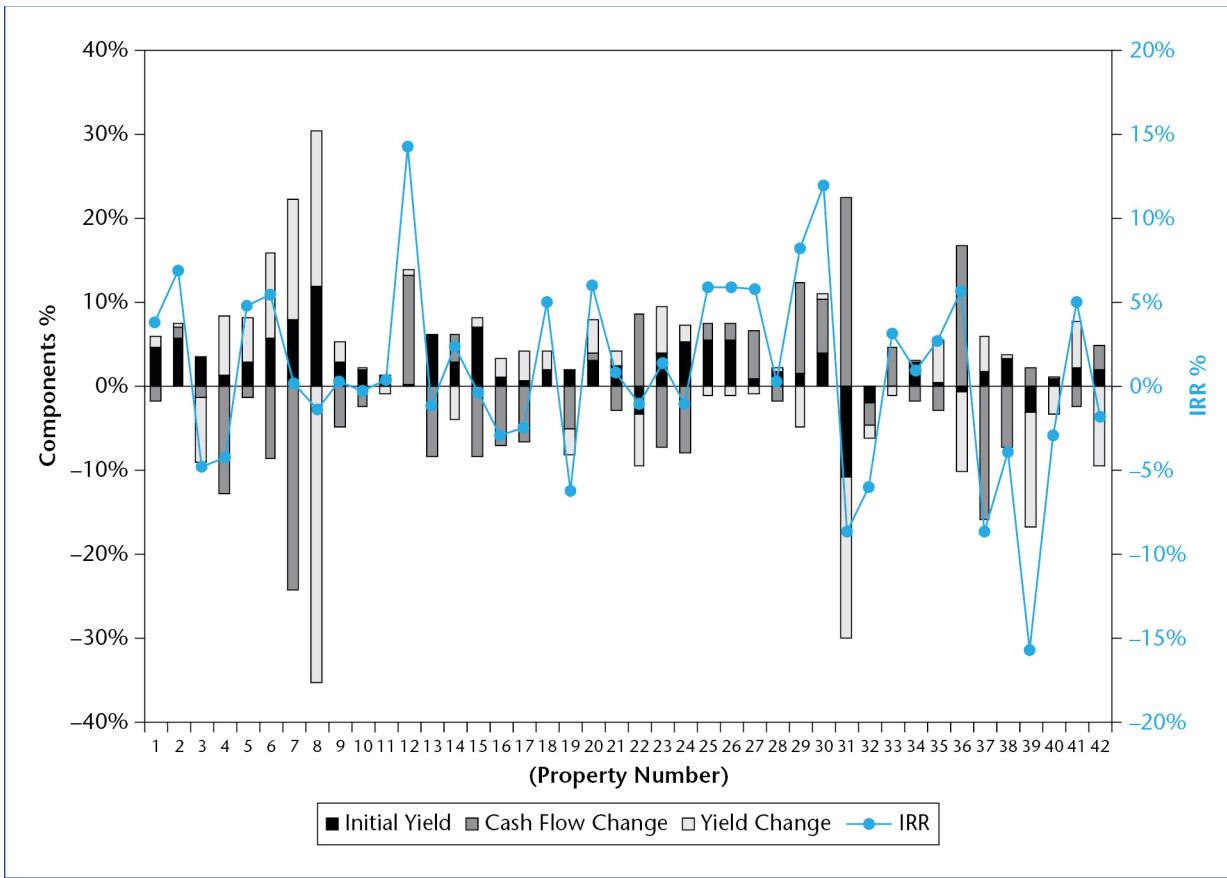


EXHIBIT 25-3 Property-Level IRR Attributes in 42 Round-Trip Investments: Subject Property vs. NCREIF Cohort, IRR Component by Property.

Source: Modified from Feng (2010).⁵

⁵ Feng, T.(2010). “Property-level performance attribution: Demonstrating a practical tool for real estate investment management diagnostics.” Master’s Thesis, MIT Center for Real Estate.

by converting the appropriate subindex of the NPI into an equivalent multiyear IRR performance. (This information is available on the NCREIF website.) Once a benchmark is selected, the difference between the subject property or portfolio performance and the benchmark is computed over the same historical time span, both overall and for each of the three performance attributes. The result will be a comparison that might look something like the bars in Exhibit 25-3, which depicts an example real-world application of the since-acquisition IRR attribution benchmarking procedure.

The exhibit shows the lifetime performances (from acquisition through disposition) of 42 randomly-numbered institutional real estate investment properties. Each property's overall IRR, as well as its three performance attributes of IY, CFC, and YC, are measured *relative to* a NCREIF benchmark. Only the difference between the property and the benchmark is shown in the exhibit (overall IRR on the right-hand vertical axis, each attribute as indicated by the bars referenced to the left-hand axis). The NCREIF benchmarks are customized to each property. The benchmark is the IRR of a NCREIF subindex representing the same geographic region and property type as the subject property and measured over the same historical holding period as the property. The 42 investments spanned the period from 1981 to 2008, with an average holding period of 10.3 years and an average property sale price of \$51 million.⁶

The first thing to notice in Exhibit 25-3 is the wide dispersion in the properties' overall IRR results around their benchmarks. This is seen in the blue dots, whose values on the right-hand axis show how the properties' lifetime IRRs differed from their NCREIF benchmarks over the same historical period. The dots are widely dispersed both above and below the horizontal axis, which represents performance just equal to the benchmark. Twenty-four of the investments beat their benchmarks while 18 fell short, with an average IRR difference of just 0.69 percent above the benchmark, implying for the average hold a cumulative outperformance of $1.0069^{10.3} - 1 = 7.3\%$. But the standard deviation of the differential performance across the 42 properties was 5.64 percent, which for the average hold would imply a one-standard deviation cumulative dispersion between +76 percent down to -45 percent.⁷

The performance attribution represented by the bars in the exhibit allows us to gain some insight into the determinants of this range of performance. The initial yield performance is relatively stable and favorable across the 42 investments.⁸ Most of the "action" in these properties' performance relative to their benchmarks lies in the other two attributes, CFC and YC, attributes that depend substantially on management of the property investments *subsequent* to the initial selection and acquisition of the properties. Close inspection reveals that it is especially the CFC attribute which, more than any other, determines whether a property beats its benchmark or falls short in terms of the overall IRR performance that is the main investment bottom line. Among the 42 properties, 57 percent beat their benchmarks, but among properties with positive CFC components (relative to the benchmark CFC), 69 percent beat their overall IRR benchmarks. All of the sample's top nine performers relative to the benchmark (in terms of IRR) had positive CFC components relative to the benchmark, while seven of the nine worst performers had negative CFC components relative to the benchmark.

⁶ The properties in Exhibit 25-3 represented all regions of the U.S. and the four core usage types of apartments, industrial, office, and retail. The holding periods ranged from 2.9 to 24.5 years, and gross IRRs (per annum, before subtracting benchmark IRRs) ranged from 29.1% (with a 5-year hold) to -1.7% (with a 4-year hold). All the properties were held as core real estate investments by a group of relatively conservative tax-exempt institutions with similar overall investment objectives and policies. Exhibit 25-3 is based on Feng (2010).

⁷ $1.0564^{10.3} - 1 = 1.76$. $(1 - 0.0564)^{10.3} - 1 = 0.55 = 1 - 0.45$. As suggested in Exhibit 25-1, this magnitude of range in individual property lifetime investment performance dispersion seems to be typical among U.S. institutional commercial properties. Similar rates of cumulative performance dispersion appear in the residuals of the repeat-sales regressions underlying indices such as the Moody's/RCA CPPI described in the preceding chapter.

⁸ Compared to the overall IRR average relative to benchmark of 0.69 percent with standard deviation 5.64 percent, the IY component had average 2.25 percent with standard deviation 3.54 percent across the 42 properties.

These results may reflect property-level investment management strengths and weaknesses of the managers of the sample properties. Relating the performance in the three attributes to the four investment management responsibilities listed in Exhibit 25-2, we can form some hypotheses. The relatively strong and reliable IY performance suggests that these managers were generally able to acquire properties more effectively than the NCREIF peer universe. Exhibit 25-2 suggests that this may have reflected some combination of good property selection and effective acquisition execution. For example, the managers may have been able to effectively utilize deep cash pockets, perhaps focused on a few target markets where they could be major players, so as to attract the best brokers to bring them privileged access to the best deals. On the other hand, once acquired, the properties in general seem not to have been as effectively managed as their benchmarks, particularly regarding ongoing cash flow during the holding period. The managers could perhaps improve their overall investment performance by figuring out how they are falling short of their peers in the management of property-level revenues, expenses, or capital improvement expenditures.

This example shows the power and potential of quantitative property-level investment analysis including performance attribution and benchmarking. However, all such formal quantitative exercises must be taken with a grain of salt and applied in an “investigative” spirit. They often bring insights by the questions they raise rather than by providing any definitive answers. They get property and asset managers thinking about how they did and what they did in specific cases, and what impact their actions had on the overall performance. These types of results should never be used mechanistically, without broader and more qualitative analysis, and they should generally be used more for constructive *diagnostic* (i.e., “learning”) purposes than for *evaluative* purposes (in our opinion).

In applying performance attribution at any level, we should also remember that return parsing is an inexact science that is sensitive to the length of time over which the return is measured. It is also important to carefully consider the meaning of property-level performance attribution measures. For example, if the yield or income component of the multiperiod return does not correspond to the yield at the time when the property was acquired, then it is not clear what implication this component of the return can logically hold regarding the investment manager’s performance, even in comparison with a benchmark. Similar questions about meaning may be raised about other components of property-level performance attribution.

25.1.3 PORTFOLIO-LEVEL PERFORMANCE ATTRIBUTION

(See 9-minute Video 2 in Chapter 25 Supplements folder.)

We now move from property-level to portfolio-level attribution. At the **portfolio level**, the key components of investment performance attribution are traditionally defined as **allocation** and **selection**. These are the two major “jobs” that an active portfolio manager may potentially perform. Such a manager may decide how to allocate the total investment capital across a range of exhaustive and mutually exclusive **segments**, that is, classes or types of assets within her purview, and/or she may select which particular individual assets to buy (and sell) within each segment. It is in the selection component that portfolio-level attribution can link down to property-level attribution. We saw in the previous section how asset selection is one of the four fundamental property-level investment responsibilities. But asset selection can also be viewed within the broader or higher-level context of portfolio management.

The terms “allocation” and “selection,” like the basic idea of performance attribution at the portfolio level, have their origins in the stock-market-based branch of the investment industry. We need to see how this tool translates into application to investment in the private real estate market. Let’s begin with a consideration of the two basic attributes: allocation and selection.

As noted, allocation refers to the chosen segment weights in the portfolio. As an attribute of investment manager performance, allocation is a very similar concept in real estate and in the stock market. In the stock market, segments may refer to industry groupings or sectors in the stock market, or to attributes of the stock, such as market capitalization or current yield. In the private real estate asset market, segments

are typically based on the space market to some extent and are usually defined by building usage type and the geographical location of properties. For example, an allocation decision would be to invest 60 percent in warehouses and only 40 percent in office buildings.⁹ The basic idea is the same as in the stock market. Segment allocation is a broader-level, more “macro” decision than selection.

Selection, on the other hand, refers to the more micro-level decision of which particular individual assets are picked by the manager to include in the portfolio. In the stock market, investment managers have more freedom of action in asset selection than in the private property market. Within a given segment, an equities manager can invest any amount in any individual asset (that is, any individual firm’s stock). In real estate, an investment manager can decide which individual properties to buy or sell. However, in real estate, only assets that are for sale can be bought; so, the manager cannot choose from among the entire potential universe of individual assets. Transaction costs are high in the property market, which usually means long holding periods must be planned for. Finally, as was the subject of the previous section on property-level performance attribution, with private real estate investment, the property owner is responsible for the operational management of the assets he selects. Indeed, as we have seen, a large part of the total investment performance may be attributable to operational management.

For all of these reasons (but especially the last), the asset selection function in private real estate is more complicated than it is in stock market investment management. At the portfolio level, it is really not possible to distinguish the component of overall performance due to asset selection per se (as understood in the stock market) from the performance component due to operational management of the assets (though this is where the property-level performance attribution discussed in the previous section can come into play). Nevertheless, there is still some interest and potential usefulness in quantitatively attributing performance at the portfolio level to allocation and selection.

To see how this is done, let’s consider a simplified example. Suppose Bob and Sue were each hired at the same time to place \$100 million worth of capital into private real estate investments in the industrial and office property segments (and then to manage those investments). Their clients gave them discretion over how much to allocate to each of the two segments. Bob immediately went out and bought \$90 million worth of industrial buildings and \$10 million worth of offices. Sue did just the opposite, buying \$10 million of industrial and \$90 million of offices. A few years later, these two intrepid investment managers got to talking at a cocktail reception at a meeting of the Pension Real Estate Association (PREA). Realizing the amazing coincidence of both being hired at the same time for the same type of investment management job, and noting that they had both placed rather strong, and opposite, segment bets, they decided to compare their performances.

The table in Exhibit 25-4 shows the returns both managers earned for their clients, both in total and broken down by segment. Much to his chagrin, Bob had to admit that Sue had beaten him. Over the three years, her annual average return was 9.7 percent, versus Bob’s 9.2 percent, a differential of 50 basis points. However, a more careful examination of the returns each manager earned by segment reveals differing sources for their relative performances. Bob clearly beat Sue in property selection performance because *within* each segment Bob’s properties earned higher returns than Sue’s. His industrial properties returned 9 percent to her 7 percent, while his office properties returned 11 percent to her 10 percent. But Sue beat Bob in allocation performance. Her decision to allocate a greater share of her capital to office than industrial properties turned out to have been advantageous because of the generally better performance of office properties as compared to industrial properties. In this case, Sue’s superior segment bet dominated over Bob’s superior selection performance to give Sue the overall advantage.

The idea in portfolio-level performance attribution is to break down the total return differential between two managers’ performances (or, more often between a manager and a relevant benchmark)

⁹ Or to invest 60 percent east of the Mississippi and only 40 percent west of the Mississippi. Or to invest 36 percent in warehouses east of the Mississippi, 24 percent each in offices east and warehouses west, and 16 percent in offices west of the Mississippi. You get the idea.

EXHIBIT 25-4

Returns Realized for Clients by Bob and Sue

Weights	Bob	Sue
Industrial	90%	10%
Office	10%	90%
Returns	Bob	Sue
Total portfolio	9.20%	9.70%
Industrial properties	9.00%	7.00%
Office properties	11.00%	10.00%

into two components, one based on allocation and the other on selection. However, the overall performance differential is a nonlinear function of the pure allocation and pure selection performance differentials, so there is no unambiguous way to attribute the total differential completely and exactly to these two sources. Nevertheless, some reasonable quantification can be attempted. To see how, let's carry on with our Sue-versus-Bob comparison.

Consider allocation performance first. Suppose we want to quantify how much of Sue's overall 50-basis-point differential over Bob is attributable to her allocation difference with respect to Bob. One way to do this would be to multiply Sue's return performance within each segment by the difference between her allocation and Bob's allocation, and add these products across all the segments. This is demonstrated in formula (1a):

$$\begin{aligned}
 A_S - A_B &= r_{SI}(w_{SI} - w_{BI}) + r_{SO}(w_{SO} - w_{BO}) \\
 &= 7\%(0.1 - 0.9) + 10\%(0.9 - 0.1) \\
 &= -5.6\% + 8\% = +2.4\%
 \end{aligned} \tag{1a}$$

The implication is that Sue outperformed Bob by 240 basis points on the basis of her superior allocation.

The fact that Sue's overall performance differential was less than this, only 50 basis points, reflects the fact that she lost ground relative to Bob in the other attribute of performance, namely, selection. In fact, the simplistic implication is that she lost 190 basis points compared to Bob as a result of her relatively inferior selection: $+2.4\% - 1.9\% = +0.5\%$.

Now suppose we apply the same reasoning to quantify Sue's selection performance differential directly. That is, we multiply the difference between her return and Bob's return within each segment by Sue's allocation weights and sum these products across the segments, as here:

$$\begin{aligned}
 S_S - S_B &= w_{SI}(r_{SI} - r_{BI}) + w_{SO}(r_{SO} - r_{BO}) \\
 &= 0.1(7\% - 9\%) + 0.9(10\% - 11\%) \\
 &= -0.2\% + 0.9\% = -1.1\%
 \end{aligned} \tag{1b}$$

By this way of figuring, Sue's inferior selection performance cost her only 110 basis points relative to Bob, instead of the 190-basis-point loss we previously computed as a result of formula (1a). If selection's effect in the differential is only 110 basis points, then allocation's contribution is not 240 basis points, but only 160 basis points, as $1.60\% - 1.1\% = +0.5\%$.

We used a consistent approach across both (1a) and (1b), basing our computation in both cases on Sue's results in the *other* attribute.¹⁰ The only difference in arriving at our two conflicting answers for the performance attribution is the order in which we did the computations. If we compute the allocation effect first, we get an implied attribution of +2.4 percent to allocation and -1.9 percent to selection. If we compute the selection effect first, we get an implied attribution of +1.6 percent to allocation and -1.1 percent to selection. Which is the real answer?

Well, how about neither. When we use either of these two approaches, we are contaminating the "pure" allocation or selection effects by arbitrarily including the effect of interaction between the two effects with one or the other. For example, in computing Sue's -110-basis-point selection differential in formula (1b), we used Sue's allocation weights. Suppose we had used Bob's allocation weights instead. Then we would have gotten the -190-basis-point effect we computed from (1a):

$$\begin{aligned} S_S - S_B &= w_{BI} (r_{SI} - r_{BI}) + w_{BO} (r_{SO} - r_{BO}) \\ &= 0.9(7\% - 9\%) + 0.1(10\% - 11\%) \\ &= -1.8\% + 0.1\% = -1.9\% \end{aligned}$$

But why should either Sue's or Bob's allocation weights be "correct" for quantifying Sue's selection effect? Sue's weights are faulty because their use combines the effect of her allocation decision with that of her property selection performance, destroying the "purity" of the measurement of her selection performance. Bob's allocation weights are similarly faulty in this context because there is no logical reason to use Bob's allocation weights to quantify Sue's selection performance. What are we to do?

25.1.4 USE OF A BENCHMARK IN PERFORMANCE ATTRIBUTION

The only way out of our dilemma is to try to identify a reasonable **benchmark**, that is, a reference point that makes sense as a basis on which to quantify a manager's performance. We will discuss the question of an appropriate definition of benchmarks for private real estate performance in the next section. For now, let's suppose that Bob is a good benchmark for quantifying Sue's performance. This might be the case, for example, if Bob's allocation weights across segments and his return performance within each segment were typical of the investment managers in Sue's peer group of investment managers with similar specialties, foci, and objectives to Sue's. Or it might be the case if Bob's segment weights and in-segment returns were broadly representative of the entire market of all the industrial and office properties in the country. In any case, let's assume that Bob's performance is a good benchmark for Sue's performance.

In this case, it makes sense to quantify the pure effect of Sue's allocation performance by computing the effect of the difference between her and her benchmark's segment weights measured *conditional on the benchmark's in-segment return performance*. Similarly, it makes sense to quantify the pure effect of Sue's selection performance by computing the effect of the difference between her and her benchmark returns within each segment, *conditional on the benchmark's segment allocation weights*. Using Bob's results as the benchmark, this perspective reveals that Sue's pure allocation performance effect is +160 basis points, and her pure selection performance effect is -190 basis points:

$$\begin{aligned} A_S - A_B &= r_{BI} (w_{SI} - w_{BI}) + r_{BO} (w_{SO} - w_{BO}) \\ &= 9\%(0.1 - 0.9) + 11\%(0.9 - 0.1) \\ &= -7.2\% + 8.8\% = +1.6\% \end{aligned} \tag{2a}$$

¹⁰ That is to say, to quantify the effect of Sue's allocation, we used Sue's selection performance, and to quantify the effect of Sue's selection, we used Sue's allocation.

$$\begin{aligned}
 S_S - S_B &= w_{BI}(r_{SI} - r_{BI}) + w_{BO}(r_{SO} - r_{BO}) \\
 &= 0.9(7\% - 9\%) + 0.1(10\% - 11\%) \\
 &= -1.8\% - 0.1\% = -1.9\%
 \end{aligned} \tag{2b}$$

Alas, these two pure effects do not add up to the total performance differential between Sue and her benchmark: $+1.6\% - 1.9\% = -0.3\% = +0.5\% = 9.7\% - 9.2\%$. This is a consequence of the non-linearity of the total return performance differential as a function of the allocation and selection effects.¹¹ As a result, the total differential between Sue's performance and her benchmark consists not only of the two "pure" effects of allocation and selection, but also of an **interaction effect** that combines these two. In fact, at a fundamental conceptual level, we cannot really account for the interaction effect in a way that is not arbitrary in terms of attributing this effect either to allocation or selection. The most logically sound approach is to leave the interaction effect separate, as a third attribute of the overall performance differential.¹² In our present example, the interaction effect in the differential of Sue's return over her benchmark is +80 basis points, as $+1.6\% - 1.9\% + 0.8\% = +0.5\%$.

The performance attributions that result from all three of the approaches we have described are summarized in Exhibit 25-5 for the numerical example we have been working with here. The first two columns in the table depict the approach based on formulas (1a) and (1b), respectively, in which results are conditional on the order in which the effects are computed. The third column depicts the approach represented by formula (2), which is independent of the order of computation and provides a logical quantification of the "pure" functional effects, with the interaction effect quantified separately.¹³

¹¹ $r_S - r_B = (w_{SI}r_{SI} + w_{SO}r_{SO}) - (w_{BI}r_{BI} + w_{BO}r_{BO})$
 $= (w_{SI} - w_{BI})r_{BI} + (w_{SO} - w_{BO}) = \text{Pure allocation}$
 $+ w_{BI}(r_{SI} - r_{BI}) + w_{BO}(r_{SO} - r_{BO}) = \text{Pure selection}$
 $+ [(w_{BI}r_{BI} + w_{BO}r_{BO}) + (w_{SI}r_{SI} + w_{SO}r_{SO})$
 $- (w_{BI}r_{SI} + w_{BO}r_{SO} + w_{SI}r_{BI} + w_{SO}r_{BO})] = \text{Interaction effects}$

¹² The interaction effect is sometimes referred to as the cross-product, as that is what it is from a mathematical perspective. Some analysts suggest that the interaction effect itself is indicative of a specific ability on the part of the manager, namely, the manager's success in specialization that leads the manager to overweight segments within which the manager has relatively superior selection skills (even though such segments might not be strategically superior from an overall allocation perspective). But perhaps this is trying to read too much from the "entrails." When all is said and done, the interaction effect is just that, the combined (multiplicative) effect of both allocation and selection.

¹³ It should be noted that the portfolio-level attribution measures described here are essentially single-period measures. In applying such attribution analysis to multiperiod spans of time, a certain amount of arbitrariness or obfuscation tends to occur because in reality segment weights will tend to change gradually over time, both in the manager's and the benchmark portfolio. In practice, alternative approaches to quantifying the attribution measures over multiperiod spans of time include (1) using the initial, terminal, or average (across the overall time span) allocation weights applied to the time-weighted average multiperiod returns within each segment; and (2) computing the single-period attribution measures within each period, and then taking the arithmetic average of the attribution measures across the overall time span. The former approach is somewhat arbitrary and inaccurate in the segment weightings. The latter approach relates to the arithmetic mean time-weighted return, not the geometric mean that is typically used for the overall multiperiod return measure in performance evaluation. However, the latter approach has a better statistical justification, as it does not violate the single-period nature of the attribution measures, but simply views each periodic return as an equally representative sample of time for indicating the manager's performance attribution. The arithmetic time-weighted mean also has the advantage of being completely decomposable into the sum of the arithmetic time-weighted average income return component plus the arithmetic time-weighted average appreciation return component.

EXHIBIT 25-5

Sue's Performance Attribution Relative to Her Benchmark

Attribute	Conditional on Manager Selection	Conditional on Manager Allocation	Unconditional
Allocation	2.40%	1.60%	1.60%
Selection	-1.90%	-1.10%	-1.90%
Interaction			0.80%
Total differential	0.50%	0.50%	0.50%

25.1.5 THE CASE FOR USING MANAGER ALLOCATION WEIGHTS IN THE BENCHMARK

Let's return for a moment to the approach represented by Formula (1b) in which the manager's allocation weights are applied to the differential between the manager's and benchmark's selection performance in order to quantify the manager's selection performance with respect to the benchmark. In fact, this approach has been employed widely in practice to quantify the manager's selection performance.

The intuitive appeal of this approach is that it can eliminate that pesky interaction effect, and it seems only fair to give credit to managers for their own allocation weights when quantifying the effect of their differential return performances within segments. But keep in mind that in so doing, one is mixing the effects of two performance functions, allocation and selection. By the same reasoning, it would be only fair to apply the Formula (1a) approach to quantifying the allocation effect (using managers' own within-segment returns to quantify the impact of their segment weight differences versus the benchmark). Of course, Formulas (1a) and (1b) do not add up to the total return differential: $2.4\% - 1.1\% = 1.3\% = 0.5\%$.

In some circumstances, however, the argument for the use of Formula (1b) to quantify the selection effect makes sense. Suppose the manager does not have control over her own allocation (e.g., the client did not give the manager discretion over the segment weights). In such circumstances, it makes sense to "customize" the manager's benchmark by recasting the benchmark segment weights to equal the manager's segment weights. In other words, the manager's benchmark should be computed with the benchmark in-segment returns but the manager's segment weights. In this case, a direct comparison of the manager's return with such a customized benchmark return will be equivalent to Formula (1b), and the entire overall performance return differential between the manager and the benchmark will consist purely of the selection effect (the pure allocation and interaction effects will be zero, by construction).

25.2 INVESTMENT PERFORMANCE EVALUATION AND BENCHMARKING

Our discussion of manager benchmarks in the previous section is a good stepping stone to our next topic, **performance evaluation**. Performance evaluation is one of the most basic and important tasks in the investment industry. In the context of investment management, performance evaluation refers to the need to arrive at some sort of judgment about how well a given manager has performed, either in absolute or in relative terms. While performance evaluation broadly involves considerations that cannot be well quantified, some aspects of the job can be quantified, and it is to quantitative evaluation that we turn our attention now.¹⁴

¹⁴ Please see the two appendices to this chapter for more depth and detail.

25.2.1 THE BASIC IDEA

A typical performance evaluation of an investment manager or an investment fund is presented in Exhibit 25-6, which depicts the evaluation of two managers, A and B. After a period of time, typically three to five years, the average return achieved over that period is computed for each manager.¹⁵ The returns achieved over that same period of time are also computed for each other manager in a **peer universe** of all the competing managers (or funds) who also specialize in the same type of investments. In Exhibit 25-6, the manager's performance is indicated by the cross. The vertical line indicates the 5th and 95th percentile range of the manager's peer universe. The box indicates the middle two quartiles, and the horizontal line inside the box indicates the median manager's performance within the peer universe. The triangle indicates the return recorded by a passive index that is broadly representative of the type of assets the manager specializes in. (A **passive index** is one whose allocations across assets remain relatively constant.)

In the example shown in Exhibit 25-6, manager A did quite well, performing in the upper quartile of her peer universe and substantially beating the relevant passive index. She will likely be rehired and may have earned an incentive fee. It doesn't look so good for manager B. The investor likely will fire manager B or reduce his allocation in the future, probably hiring one of his competitors instead. (Note that manager A probably is not a competitor to manager B, as they are in two different asset classes or management styles.¹⁶)

25.2.2 BENCHMARKS AND INVESTMENT STYLES IN THE PRIVATE REAL ESTATE ASSET CLASS

In Exhibit 25-6, the manager's performance is compared both with respect to the median of her peer universe and with respect to a passive index broadly representative of her asset class. For example, in the public equity investment business, if the manager specializes in large-capitalization stocks, then the peer universe would consist of all the major mutual funds with a similar specialty. The most widely used passive index to benchmark large-stock performance is the Standard & Poor's Composite 500 Index (the S&P 500).

In the example shown in Exhibit 25-6, manager A has beaten both the benchmark index and her average competitor. Manager B has been beaten by both of these reference points. More generally,

¹⁵ Normally, the time-weighted average return is used (in particular, the geometric average) because at the portfolio or fund level (as distinct from the property-level), the investment manager does not usually have much control over the timing of capital flow. As noted in Chapter 3, time-weighted average returns are neutral with respect to capital flow timing, while dollar-weighted returns reflect the effect of such timing. If the manager does substantially control the timing of capital flow, then the IRR might be a more appropriate measure of the manager's performance. However, the benchmark is meant to be neutral with regard to performance actions. In the case of capital flow timing, such neutrality is represented by the time-weighted average return. The geometric mean is typically used rather than the arithmetic mean because performance evaluation is a backward look at history (rather than a forecasting exercise in which the statistical properties of the arithmetic mean would be useful), and because the geometric mean is not sensitive to the volatility in the periodic returns. (The arithmetic mean is greater the more volatile the returns.) It should also be noted that at the portfolio or fund level being addressed here (as distinct from the property-level considered in Section 25.1.1), the periodic HPRs are normally computed as value-weighted (cross-sectional average) returns within each period. That is, within each period, the return to each property is not weighted equally but by the relative value of the property within the portfolio. This type of value-weighting is necessary to enable the HPR each period to reflect the actual return to the portfolio within that period. The focus here (and in the portfolio-level attribution analysis described in Sections 25.1.2–4) is on portfolio-level returns, rather than viewing each property as a sample of the manager's performance.

¹⁶ The placement of manager A's and manager B's performance evaluations together in Exhibit 25-6 is done merely for illustrative purposes. It is not meant to imply that there should be a head-to-head competition between these two managers. Such direct comparisons would normally only take place between managers within the same asset class or segment (or "style") of investments, to avoid an apples-to-oranges comparison.

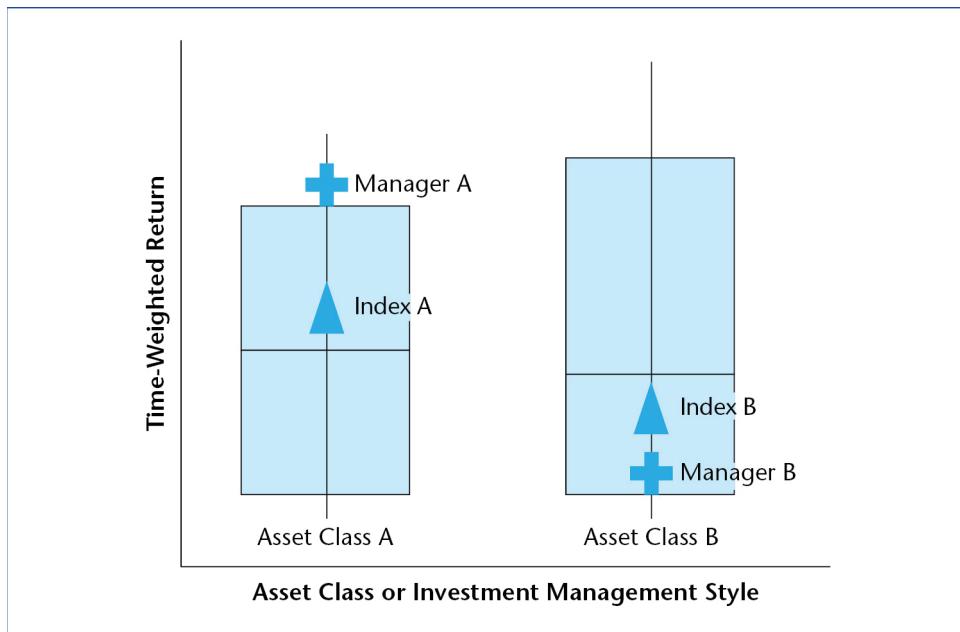


EXHIBIT 25-6 Graphical Depiction of Typical Investment Manager Performance Evaluation.

however, there is potential ambiguity in this process. A manager might, for example, beat the average competitor but fall short of the relevant passive index, or vice versa.

So, which type of reference point is better as a benchmark for judging manager performance? Should it be some central tendency of the manager's peer universe (say, the mean or median), or should it be a passive index? In the securities investment industry, both types of benchmarks are used. But in private real estate a truly passive index is not possible. Fundamentally, this is because of a characteristic of the property market which has been prominent in much of our previous discussion, namely, private real estate investment managers have to be responsible for the operational management of the properties they own on behalf of their clients. Furthermore, as we have noted, private real estate investment necessarily involves long holding periods because transaction costs are high, leading to our previously highlighted fact that much of the overall investment performance is attributable to how well properties are managed while they are held. This is **active management**, of necessity.¹⁷ Passive indices make sense as benchmarks if they represent a plausible alternative that the investor could actually invest in, as through an index fund. But, at least unless or until index derivatives (often attempted, with limited success so far) are available, there can be no such thing as a passive index or an index fund in private real estate.

¹⁷ The term *active management* in the investment management industry has traditionally been defined from a public securities market perspective. In that context, of course, active management does not imply operational management of the investment assets, as this is not possible for investment managers in the securities industry. But in a broader sense, active management refers to the attempt to beat the market, or to outperform the average or typical portfolio of assets, by whatever means are at the disposal of the manager. In the case of securities investment managers, this is limited to the use of investment research and the employment of active trading strategies or target allocations. Private real estate investment managers may employ these types of activities, broadly defined, although in a more sluggish asset market, dealing with assets that are primarily income generators rather than growth plays, and in an environment of high transaction costs. But in addition to these traditional active management tools, private real estate investors can, indeed must, use their control over property-level operational management of the assets to try to improve their investment performance.

Instead of passive indices such as the S&P 500, the only type of periodic return indices available to serve as benchmarks in the private real estate asset class are peer group indices. These are indices, like the NCREIF and MSCI/IPD indices, that essentially reflect the performance of all (or most) of the private real estate investment managers in the country or in the relevant branch of the investment industry. As such, peer group indices correspond conceptually to the central tendency of a given manager's relevant peer group. They make sense as benchmarks because in some sense they represent the average of the "alternatives" that the investor could have achieved if she hadn't picked the particular investment manager that she hired.

The traditional and most widely cited NCREIF and MSCI/IPD indices are property-level indices, which track the performance of the underlying physical assets. Property-level indices were our focus in Chapter 24. Such indices can be useful as benchmarks for property-level performance analysis and evaluation, as described in Section 25.1. However, passive investors who make their real estate investments through investment managers actually experience not the underlying property performance directly, but that performance as managed and realized and passed on to the investor by funds or investment vehicles set up by the investment managers. It is the performance of such funds that investors need to benchmark when evaluating the performance of their hired managers. Both NCREIF and MSCI/IPD produce so-called **fund-level indices** for the U.S. institutional real estate investment industry. These are indices that attempt to track the performance actually realized by investors in funds that invest in properties, rather than the performance achieved directly by the underlying properties.

Exhibit 25-7 presents the cumulative total returns (income plus capital appreciation) from 2000 through late 2023 as tracked by the NCREIF Property Index (NPI) and by two institutional investor fund-level indices, again published by NCREIF. The first fund-level indices represent all funds and the second zooms in on the "core" investment **style**. An investment style in this context refers to the types of assets that the investment manager invests in, as well as certain key aspects of the management policies, especially the amount of financial leverage that may be used by the manager. The concept of investment style is a bit fuzzy, as is the exact definition of any given style, but it is important. We will have more to say about styles shortly, but for now let's consider the two fund-level indices in Exhibit 25-7.

The black line in Exhibit 25-7 tracks the performance of investment funds self-described as having a **core** style of investment. This means that they focus on relatively safe fully stabilized (largely leased-up) properties, using relatively little leverage in the investments. Core investment funds have relatively low target returns, but they are also supposed to have little risk. In terms of physical assets, classical institutional quality office, apartment, retail and industrial properties of the type that largely make up the NPI would be typical investments of a core style fund. Traditionally, core investment funds employed little or no financial leverage, but by the 2000s, they were often using modest amounts of debt, up to 20 percent to 30 percent loan-to-value ratios (LTVs).

Based on value-weighted all-funds indices, net of fees. NPI is the property level gross of fees. The black line above in the exhibit tracks the performance of the NCREIF All Funds Index that encompasses a broad spectrum of open-end, commingled equity real estate funds with diverse investment strategies. These funds vary in sector and regional focus, leverage levels, property life cycles, and include both core and non-core investment styles and capture a wide array of investment styles, ranging from core to non-core strategies, reflecting the diverse approaches of open-end real estate funds in the market. Consequently, the All Funds Index is not intended as a performance benchmark but serves as a market index representing real estate vehicles with liquidity provisions. In addition to core funds, it covers investment styles that are more aggressive: Value-added funds, for instance may invest in assets that are not fully stabilized (still need some leasing) or in need of some turnaround management. Or they may be smaller or less "prime" properties or locations. Perhaps most significantly, the value-added style traditionally allows considerably more financial leverage than the core style, with typical LTVs in the neighborhood of 50 percent or slightly more. This is still conservative

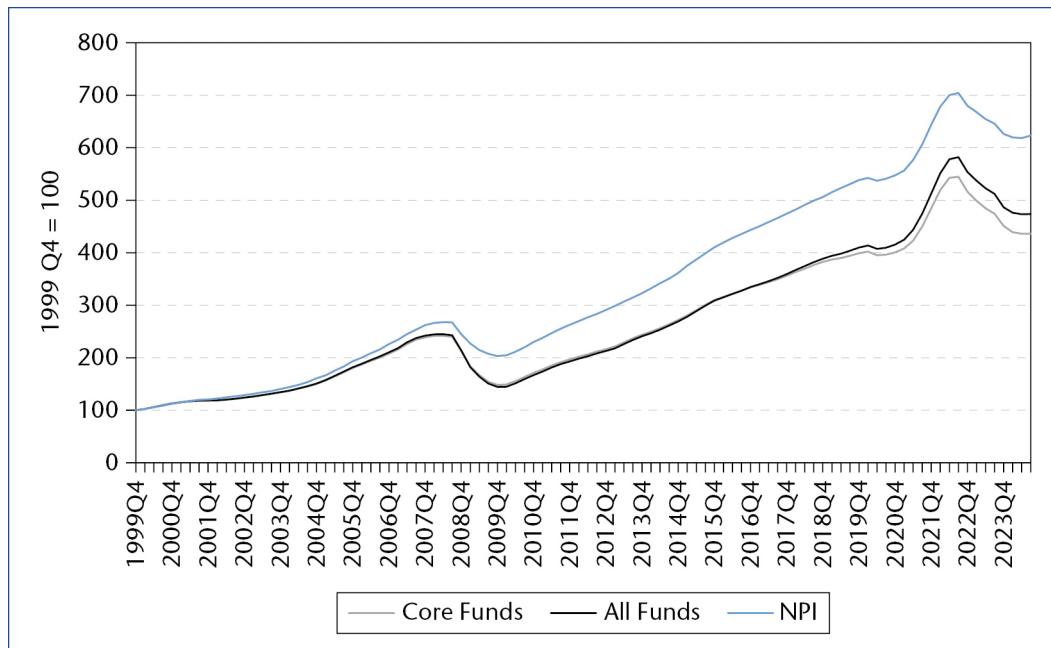


EXHIBIT 25-7 NCREIF Fund Indices vs. NCREIF Property Index: Cumulative Total Returns, 2000–2023.

by the standards of many real estate investors, but as you recall from Chapters 12 and 21, the impact of leverage on investment risk is nonlinear in the LTV ratio.¹⁸

Since core funds invest in largely the same types of properties that compose the NPI and with little leverage, the performance of the core funds index in Exhibit 25-7 (gray line) can be compared to (and is largely based on) the property-level performance represented by the blue line tracing the NPI. We see that during 2000–2023, the core funds on average provided less total return to their investors than the underlying properties in the NPI. The obvious reason for this is that the fund-level index is net of the funds' investment management fees that they charge their investors, whereas the NPI, being a property-level index, does not reflect such charges. Core funds typically charge close to 100 basis points per year to cover the expenses and necessary profit of the investment management. Another major reason why the core fund index performed below the underlying properties during 2000–2023 was the way the funds on average managed their financial leverage, as we will discuss shortly.

These differences between the core fund index and the NPI are further magnified for value-added funds. Value-added funds are more management intensive than core funds and therefore may charge higher asset management fees (a component of which may be performance based). And the greater financial leverage in the value-added funds did not work to their advantage during the global financial crisis period.

How can financial leverage cause fund performance to fall below that of the underlying physical assets the fund is investing in, when the leverage may have appeared to be "positive" at the time

¹⁸ Recall from Chapter 12 that the effect of leverage on risk is approximately characterized by the leverage ratio: $LR = 1/(1 - LTV)$. Thus a 50% LTV equity investment has twice the risk of the underlying properties. And indeed we see that the Value Added Index has about twice the quarterly volatility of the underlying property index (NPI).

when the loans were taken out?¹⁹ Let's take a simple example. Suppose in the up part of the asset price cycle property values rise from 100 to 150. Suppose by the end of that period an investment fund has a 33 percent LTV ratio, say, \$100 in equity and \$50 in debt. Now suppose property values fall by one-third, down to 100. The fund still owes \$50 and so has lost half its equity value and it now has a 50 percent LTV ratio. At that point the fund probably faces considerable pressure to delever, as it is not clear what the future will be in the property market. Somehow the fund raises \$17 in new equity, which it uses to pay its debt down to \$33 to return it to a 33 percent LTV. Then suppose the property market partially recovers, to a value of 125, leaving the fund with \$92 in equity and \$33 in debt (to total up to the \$125 asset value).

Now for simplicity, assume each of the asset price moves occurred in one period and let's compute the simple holding period returns (HPRs), the types of returns tracked by periodic return indices such as those in Exhibit 25-8. In the first crash period (2008–2009), the property assets had a -33 percent return (from 150 to 100). The fund had a -50 percent return (from 100 to 50), with its loss compared to the property market exaggerated by its 33 percent LTV at the peak. At this point new equity investors come into the fund (bringing \$17 to pay down the debt), diluting the preexisting investors. Such a recapitalization is not reflected in a periodic return index like those in the exhibit. In the first years of the recovery period, from 2010 to the end of 2012, properties gained

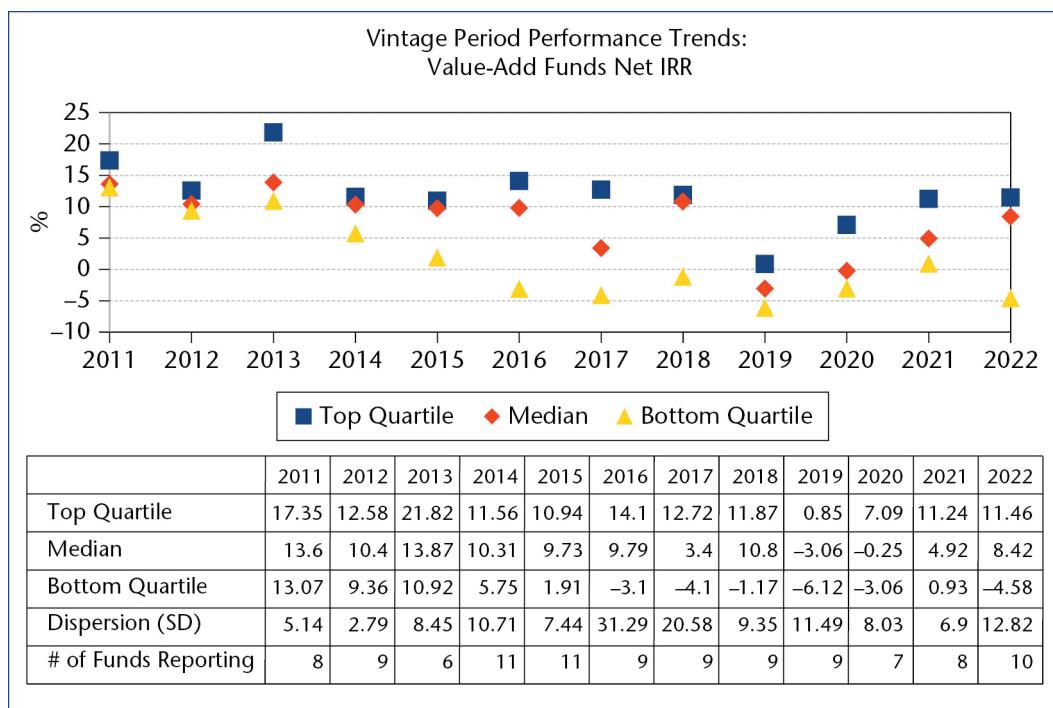


EXHIBIT 25-8 NCREIF Report of Real Estate Value-Add Funds IRR Performance, by Vintage Year (as of September 30, 2024).

Source: NCREIF Fund Index, Closed-End Value-Add, NCREIF, Chicago (2024).

¹⁹ Recall from Chapter 12 that “positive leverage” refers to the circumstance where the interest rate on the loan is less than the return on the property, implying that the leverage should increase the equity return on average. (See Section 12.4.)

25 percent (from 100 to 125). The fund gains 37 percent (from 67 to 92). Across the two periods the property-level HPRs are therefore -33 percent and +25 percent for a time-weighted arithmetic average return (TWR) of -4 percent. The fund-level HPRs are -50 percent and +37 percent for a TWR of -6.5 percent, 250 basis points lower than the underlying properties.

This simplified story is the essence of how poor timing (in retrospect) in the management of leverage during the boom-and-bust cycle of the 2000s led to widespread fund performance below that of the underlying property assets. This happened not only in private real estate investment management but in many REITs as well (and not just in real estate but in many industries in the 2000s).

TEXT BOX 25.1 LIKE FOR LIKE COMPARISON OF RISKS AND RETURNS

Traditional attribution analysis implicitly assumes that the manager's portfolio is the same risk as the benchmark when decomposing the relative or active return into selection and allocation components. If some sectors are riskier than others, or the manager has invested in riskier assets in a sector, the manager may appear to have made better selection or allocation decisions when the results were due in part to taking on more risk. Using beta as a measure of risk, Fisher and D'Alessandro (2019) show that sectors for both the manager's portfolio and the benchmark should be adjusted to the same level of risk before using the returns in attribution analysis. The manager may have outperformed the benchmark by allocating more capital to riskier sectors or investing in riskier assets within a sector or underperformed but with a less risky portfolio.²⁰

Let us return now to the real estate institutional investment management styles, and note that there is a third major style, referred to as **opportunistic**. This is the most aggressive of the three styles. Opportunistic investment managers and funds aim for high returns (15 percent to 20 percent was a typical target in the last two decades) which can only be achieved by taking on considerable risk. Opportunistic funds may invest in a broad range of asset types, including land, development projects, foreign properties (possibly in emerging markets), distressed debt, or properties or operating entities in need of considerable turnaround. As with the value-added investment style, if opportunistic investing is to actually add value, that is, to "beat the market" in the sense of generating positive-NPV investments (and thereby achieve high returns without commensurate risk), it generally must do so through the managers' capabilities with physical assets. However, again as with value-added investing, there is a temptation to reach for high returns simply by taking on more financial leverage and moving out along the Security Market Line (SML) in that way, achieving greater return expectations relatively easily with no real added value but with commensurately greater risk exposure. Opportunistic funds typically employ high amounts of financial leverage, often 50 percent to 70 percent LTVs in spite of more risky underlying assets. This temptation will tend to be greater during boom periods when debt capital is plentiful and interest rates are low.

²⁰ See "Risk Adjusted Performance Attribution" by Jeff Fisher and Joseph D'Alessandro, (2018), available at SSRN: <https://ssrn.com/abstract=3232392>, and "Risk-Adjusted Attribution Analysis of Real Estate Portfolios", by the same authors, *Journal of Portfolio Management*, 45:7 (2019): 80–94.

TEXT BOX 25.2 IS IT A SAMPLE OR IS IT A POPULATION: THE STATISTICAL NATURE OF BENCHMARK INDICES

Benchmark property-level real estate investment performance indices such as the NCREIF index are often spoken of in common parlance as being a **sample** of a larger underlying population of all the commercial properties of a certain class or type. Indeed, much basic research conducted using such indices effectively treats property indices this way, as they are taken, for example, to represent the performance of an entire asset class (which is generally somewhat vaguely defined). However, it is important to realize that, in the context of performance evaluation of investment managers, benchmark indices such as the NCREIF index are not playing the role of statistical samples of some larger underlying population of properties. Rather, such indices in this context are themselves meant to define an entire **population**, namely, the population of all properties held by the relevant peer group of managers.

In the field of statistical inference, there is a large conceptual difference between a population and a sample of a population.

One implication of this difference is that, if the index is viewed as a sample of a larger population, then, at least ideally, it should not be a value-weighted index, that is, an index constructed as the cross-sectional average of the individual property returns each period weighted by the relative magnitude of the property value within the overall population. Instead, the ideal index from a statistical sampling perspective would be equally weighted across properties within each period. However, the NCREIF index is, quite intentionally, a value-weighted index. This suggests that the designers of that index intended that it be viewed as a population rather than as a sample (NCREIF allows members to easily download custom indices that can be constructed by equal-weighting the property-level returns in the database. The transaction-based version of the NCREIF Index described in Chapter 24, the NTBI, is published in both value-weighted and equally-weighted versions).

Opportunistic and value add funds generally should not be benchmarked using the same type of investment return performance metric as the other real estate investment styles. The indices for core and all funds in Exhibit 25-7 are based on period-by-period holding period returns (HPRs) and focused on time-weighted returns (TWRs) to summarize their average performance over time. As described in Chapter 3, such returns do not reflect the timing of when investment capital is placed, or the effect of having different amounts of capital invested at different points in time. TWRs are appropriate for benchmarking investment managers who do not have discretion over equity capital flow into or out of the investments. This is true for all core funds and for some value-added funds, but not for most opportunistic funds. The latter function more like venture capital or private equity funds that set up a capital commitment from the investor and then “call” money into the investments when that is deemed necessary. The preferred return metric to summarize performance over time in such circumstances is a money-weighted average return, most typically measured by the internal rate of return, the IRR (as also defined in Chapter 3).²¹

Exhibit 25-8 is based on an example IRR-based performance report from the NCREIF Indices. When benchmarking using IRRs, it is very important to control for the “vintage,” the year or period of time when the peer universe of comparable funds was initiated. It is not fair to compare the IRRs of funds that started at substantially different points in history. The report shows the median and

²¹ Recall that there is a similar argument for using the IRR metric for the type of property-level performance attribution we described in Section 25.1.2. IRR-based metrics may also be appropriate for some value-added funds, depending on how they are structured and report to their investors.

interquartile range of the funds' reported net IRRs as of the current reporting date (September 2024 in the exhibit example). Note that the funds are grouped into 12 vintage year cohorts, one for each year from 2011 through 2022, corresponding to the years when the funds were launched. The "vintage year" is defined as the year when the fund first made its capital calls to its investors and began acquiring real estate assets. The report also shows the number of funds in each vintage and the cross sectional standard deviation among their IRRs as of the current reporting date (labeled "dispersion" in the report). The net IRRs are net of advisory fees, incentives, and promotes (performance-based fees as described in Chapter 15, Section 15.3), and are based on actual net cash flows to and from investors plus "reversion" values based on current estimated market values of the net assets remaining in the funds (typically evaluated by appraisals).

In Exhibit 25-8, note that as of 2024, most of the value-added funds originated between 2011 and 2022 had not achieved the 13–16 percent returns that were often stated as the target. Only the 2013 cohort of six funds were successful in that regard. Indeed, in only five of the 12 cohorts the median fund exceeded a 10 percent IRR, the 2011–2014 vintages launched at the beginning of the post-GFC recovery period. The worst performing cohorts were those that originated during the peak of the asset price cycle, not surprisingly, during 2019–2020. The 2019 cohort of 9 funds had a median IRR below –3 percent and a top quartile of a mere 0.85 percent.

Exhibit 25-9 summarizes the three major styles of institutional real estate investment management. As noted, the boundaries and distinctions between them are fuzzy. The most traditional and well-defined style is core, but it is also rather narrow. Especially with the yield compression of the twenty-first century (the general secular decline in real interest rates and expected returns), there has been a tendency to subtly expand the core style to try to achieve a bit higher returns (and thereby perhaps filling a sort of perceived gap between the value-added and traditional core styles). This has come to be known as "core-plus" (or "core+"), as indicated in the Exhibit. Exhibit 25-10 offers another example of the changing definition of "core" by a leading real estate investment manager (LaSalle).

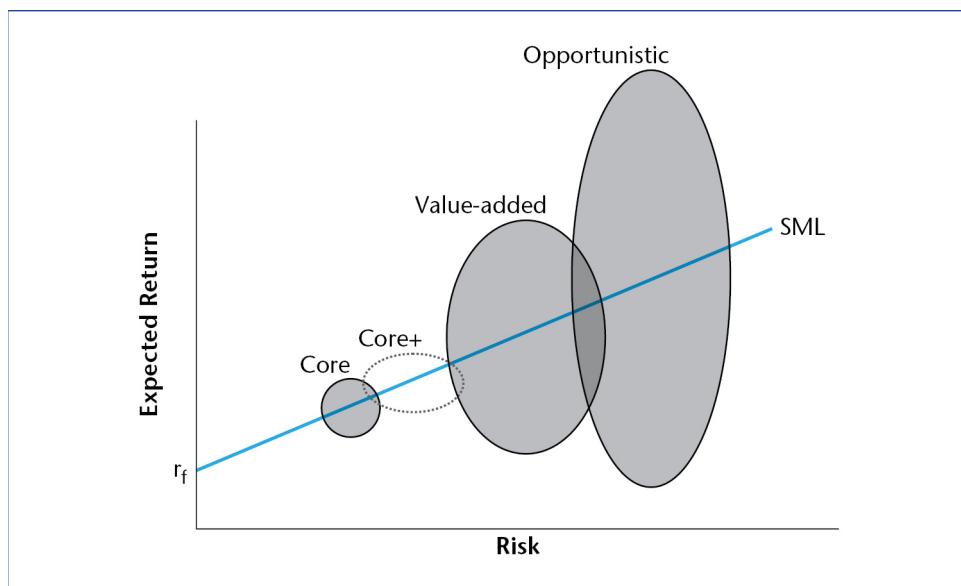


EXHIBIT 25-9 Schematic Diagram of Institutional Real Estate Investment Management Styles.

EXHIBIT 25-10

LaSalle's Changing Definition of Core

Traditional Core Mind-Set	"New" Core Mind-Set
Long Leases	Observed Long-Run Income Resilience and Growth Potential
Credit Tenants	Low Sensitivity of Cash Flows to the Economic Cycle
Lease Clauses That Pass Inflation on to Tenants	Market Conditions Pass
Minimal Near-Term Capex	Inflation Into Market Rents
Low Operational Intensity	Predictable Long-Term Capex
Gateway City Locations	Established Operating Model
Traditional Sectors: Office, Retail, Industrial, Multifamily Residential	Locations With Long-Term Liquidity and Tenant Demand
	Any Sector That Offers The Above

Source: Brian Klinksiek (2024). "The Changing Definition of Core Real Estate," *PREA Quarterly*, Spring 2024.

25.3 CHAPTER SUMMARY

With this chapter, we conclude Part VI, which focused on private investment at the macro-level of real estate investment. This chapter focused in particular on some of the basic principles, tools, and procedures used in the investment management industry, with particular reference to their use in the private property asset class. We described in some detail the techniques of investment performance attribution at both the macro property-level and the portfolio level. We pointed out the importance of formal, quantitative investment performance evaluation in real estate. Finally, we described the nature and role of such evaluation, including some subtleties and unique features that distinguish the private real estate equity asset class in this regard.

KEY TERMS

- Performance attribution
- Benchmarking
- Macro property level (performance attribution)
- Property selection
- Acquisition transaction execution
- Operational management during
- The holding period
- Disposition transaction execution
- Portfolio level (performance attribution)
- Allocation
- Selection
- Segments (in the asset market)
- Benchmark
- Interaction effect
- Performance evaluation
- Peer universe
- Passive index
- Active management

- Fund-level indices
- Investment styles
- Core
- Value-added
- Opportunistic
- Vintage
- Sample
- Population

STUDY QUESTIONS

Conceptual Questions

- 25.1. What are the three major performance attributes, and the four property-level investment management responsibilities?
- 25.2. What is the argument for basing property-level performance attribution on the since-acquisition IRR instead of time-weighted returns (TWR)?
- 25.3. What are the two major performance attributes typically identified at the portfolio management level?
- 25.4. What is meant by the term *interaction effect* in portfolio-level performance attribution? What issues arise when analysts lump the interaction effect in with either allocation or selection?
- 25.5. Piet and Yongheng are two apartment property investment managers hired one year ago by two different investors. In both cases, the managers were free to use their own judgment regarding geographical allocation between properties in the East versus the West of the country. Piet allocated his capital equally between the two regions, while Yongheng placed three-quarters of his capital in the Western region. After one year, their respective total returns were as depicted in the following table. As you can see, Piet beat Yongheng by 100 basis points in his total portfolio performance for the year.

Yongheng's and Piet's Returns Realized for Clients

Weights	Yongheng	Piet
East	25%	50%
West	75%	50%
Returns	Yongheng	Piet
Total portfolio	9.50%	10.50%
East	8.00%	9.00%
West	10.00%	12.00%

- a. How would you attribute this 100-basis-point differential between allocation and selection performance only (no other component) and you wanted to condition your attribution on computing the allocation performance component first using Piet's return performance (as in Formula (1a) in Section 25.1.3)?
- b. How would you attribute this 100-basis-point differential between allocation and selection performance only (no other component) and you wanted to condition your attribution on computing the selection performance component first using Piet's allocation weights (as in Formula (1b) in Section 25.1.3)?

c. How would you attribute this 100-basis-point differential among three components: pure allocation performance, pure selection performance, and a combined interaction effect, if you wanted to compute an *unconditional* performance attribution that was independent of the order of computation and based on Yongheng's results as a type of benchmark or standard?

25.6. In the previous question, suppose that Piet did not have the discretion to choose his allocation weights between the East and West regions, but was specifically requested by his client to place capital equally between the two regions (as he did). Suppose further that the within-region performance achieved by Yongheng represents an appropriate benchmark for property-level performance within each region. Then what is an appropriate benchmark for evaluating Piet's performance, and what is Piet's performance differential with respect to that benchmark?

25.7. What is the difference between a peer universe and a passive index, as a benchmark for evaluating investment manager performance? Why are the arguments against peer universe benchmarks less persuasive in the private real estate market than in the stock market?

25.8. In what sense (or in what role or uses) is the NCREIF index best viewed as a population (in the statistical sense)? In what sense (or in what role or uses) is the NCREIF index best viewed as a sample?

25.9. What are the two major reasons why the NCREIF-Townsend Core Fund Index exhibited greater risk (volatility and down cycle magnitude) and less total return than the underlying NPI property-level index during the 2000–2012 period, even though the core funds mostly hold the same properties as those tracked by the NPI?

25.10. Why or how did value-added and opportunistic funds perform worse than core funds during 2000–2012, even though such funds are supposed to provide higher returns over the long run than core, because they take on greater risk?

25.11. Why is it often not appropriate to track the performance or to benchmark opportunistic funds using periodic returns time series and time-weighted average holding period returns (TWRs)?

25.12. What is meant by a real estate index swap, and why is it referred to as a “futures” contract?

25.13. Describe two reasons why each party (the long position and the short position) might be willing to pay at least part of a “bid-ask spread” in a real estate index swap contract.

25.14. Suppose the risk-free interest rate is 3% and the equilibrium expected total return risk premium on property investments is 3.5% (hence, the equilibrium expected total return on real estate investment is 6.5%). What will be the equilibrium (bid-ask spread midpoint) price (fixed-leg rate) for a real estate index total return swap: (a) If the underlying index is in equilibrium (no lag effect)? and (b) If the expected average annual total return on the index during the period of the contract is 12% (possibly reflecting a lag in the index or the underlying market), but assume that the index exhibits the same risk as the property market such that 3.5% is the appropriate risk premium? (c) Same as (b) only now assume that the lagging and smoothing in the index gives it less risk than the underlying properties whose returns it measures (correctly over the long run) such that the appropriate risk premium for the index is only 2.5% (instead of the 3.5% for direct property investment); and finally (d) Suppose the conditions in (c) except that the index underlying the swap contract now is not a total return index but rather only tracks capital returns or price growth such that its expected return over the contract period is now only 5% (instead of 12%, i.e., there is expected to be 7% income return during the period).

Part VII

Other Selected Topics

Much has been covered in the book, but we have saved a few key topics for last. Earlier in the book, we spoke about a unique aspect of the real estate market: investment exposure can be obtained through the public and the private markets. So far, we have mostly addressed the latter. The first of the four chapters in Part VII, Chapter 26, delves deeply into public real estate: real estate investment trusts (REITs). Starting with an overview of the historic development of that market, both in the U.S. and across the globe, it will then move on to the institutional characteristics of REITs. The main body of the chapter looks at valuation issues, discussing the different approaches to valuing a REIT. The chapter ends with some considerations regarding REIT management and investment.

Chapter 27 expands the perspective to the global level and looks at international real estate investment. Real estate markets are very local in nature, and information at the micro location level is important for investment performance. On the other hand, capital flows into real estate are increasingly global. The question is how a far-away investor can still perform well in these local markets. The chapter also looks at the reasons for going global, as well as the obstacles to doing so.

Chapter 28 addresses an increasingly important systematic risk facing the global real estate markets: climate risk. The chapter first looks at the role of real estate in climate change, and then provides tools for real estate investors to decide whether to invest in sustainability improvements of their assets. We will then assess the different physical climate risks facing the real estate markets, the time frame over which these risks will play out, and the ways investors can deal with them.

Finally, Chapter 29 provides our vision on the impact of technology on the real estate industry and market. Technology has the potential to increase efficiency, productivity, add automation of mundane tasks, improve communication speed and the timeliness of responses to crises, and much more. Blockchain, sensors, and innovative apps are rapidly being integrated into how we buy, appraise, finance, manage, market, communicate, and even how we own real estate. These are some of the topics we will review in this final chapter of the book.



Taylor & Francis
Taylor & Francis Group
<http://taylorandfrancis.com>

26 Real Estate Investment Trusts (REITs)

CHAPTER OUTLINE

- 26.1 Introduction to listed REITs
 - 26.1.1 The Development of the Global REIT Market
 - 26.1.2 Operational Structure and Regulatory Constraints
- 26.2 REIT Analysis and Valuation
 - 26.2.1 Introduction to REIT Earnings Measures
 - 26.2.2 Valuing REITs as a Stream of Cash Flows: The Gordon Growth Model
 - 26.2.3 Valuing REITs as Collections of Assets: Parallel Asset Markets and NAV-Based Valuation
 - 26.2.4 REIT Public/Private Arbitrage
- 26.3 REIT Management Strategy
- 26.4 REIT Investor Considerations
- 26.5 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The basic regulatory constraints faced by REITs and some of the unique accounting terminology and conventions used in the analysis of REIT stocks.
- The basic process of valuation at the investment entity level, specifically for publicly listed REITs.
- The difference between growth and income stocks and the nature of REITs in this regard.
- Some of the major considerations and objectives in REIT management strategy.

Earlier chapters introduced you to the “securitization revolution” of the 1990s that led to the growth and development of the public REIT and CMBS markets and a strengthening of the link between private real estate asset markets (Main Street) and wider capital markets (Wall Street). Publicly traded equity REITs are companies that own, and in most cases actively manage, portfolios of properties.¹ A typical large REIT provides vertically integrated commercial real estate goods and services to customers who are tenants or, indirectly, other users of the built space, such as shoppers using a mall. A single firm may integrate raw land acquisition and development, portfolio management (also known as **asset management**), and operational-level property management.

¹ Not all REITs focus on property investment. Mortgage REITs invest in mortgages, either through direct lending or by purchasing mortgages and/or mortgage backed-securities. Hybrid REITs own a portfolio consisting partly of properties, and partly of mortgages. More recently, the REIT structure has been used for infrastructure assets, like cell towers. In addition, not all REITs are publicly traded. Non-publicly-traded REITs may be “incubators” that hope to go public, or they may be a convenient way to organize what is essentially a private equity real estate fund. The present chapter focuses only on publicly traded equity REITs, which is the dominant component of the global REIT industry by market capitalization and trading activity.

REIT shares are publicly traded in stock exchanges all over the world.² But REITs are *real estate* firms, **pure plays** in the sense that their assets and activities are largely restricted to real estate because of their claim to REIT tax-exempt status, which allows them to avoid corporate level income tax. Consequently, it is appropriate to view REITs as macro-level real estate investment entities.

With this in mind, the purpose of this chapter is to introduce you to the modern REIT industry in more depth than we have already covered it, and largely from a macro perspective. We will focus the second part of the chapter on the process and nature of REIT valuation in the stock market. In addition, we want to touch on some practical and strategic issues regarding the REIT industry from both REIT management and investor perspectives.

26.1 INTRODUCTION TO LISTED REITS

What exactly are REITs, why do they exist, and how and why have they become an important part of the real estate investment universe? We start by providing a brief history of the global development of the REIT structure, introducing some of the peculiar features and terminology associated with it as we go along, and we will then discuss the typical operational structure as well as the regulatory restrictions REITs face as a consequence of their favorable tax status.

26.1.1 THE DEVELOPMENT OF THE GLOBAL REIT MARKET

The REIT investment vehicle was created by the U.S. Congress in 1960 through legislation called the Real Estate Investment Trust Act, which authorized a real estate ownership structure with tax treatment similar to that of mutual funds: a pass-through entity that distributes most of its earnings and capital gains, without the obligation to pay corporate tax. The idea was to do for commercial real estate investment what mutual funds had done for stocks; to provide small individual investors (“retail investors”) a means to invest in a diversified portfolio of many individual assets without requiring a huge fortune. Listed REITs thus offer investors a *liquid* way to invest in a *diversified* portfolio of commercial property. At the same time, they provide a way for commercial property to obtain equity capital financing via the public stock market. Outside of the U.S., the Netherlands (1969) and Australia (1971) were early adopters of this idea. As of 2022, REITs exist in at least 40 countries.

Until 1986, US REITs were prohibited from being “self-administered.” That is, REITs were required to be externally managed. A separate investment management entity, in which the REIT shareholder had no stake, managed the REITs’ property holdings. The idea was that the REIT was to be a largely “passive” entity, analogous to a mutual fund that does not actively manage the firms in whose stock the mutual fund is invested. However, real estate assets are not like publicly traded stocks. Their owners must manage the properties. The third-party administration requirement greatly hampered the development of the US REIT industry.

With the removal of this restriction in the 1986 tax reform, the REIT industry gradually discovered a completely new way of operating, which ushered in the so-called “Modern REIT Era” in the 1990s. Today, REITs tend to be internally managed, in the sense that the REIT shareholders ultimately govern the firm, and the firm, in effect, directly manages the portfolio of properties held by the REIT. This reduces agency problems between REIT managers and shareholders, and it has been a major factor in making REITs the players they now are in all sectors of the institutional commercial property market.

Exhibit 26-1 illustrates this for the U.S. and shows the extent to which different property types are accessible through the REIT market at two points in time: 2000 and 2022. In 2000, the property types available through REIT investment mostly involved the ‘traditional’ commercial real estate

² We use the term “REIT” throughout this chapter, irrespective of the name this structure has in different countries.

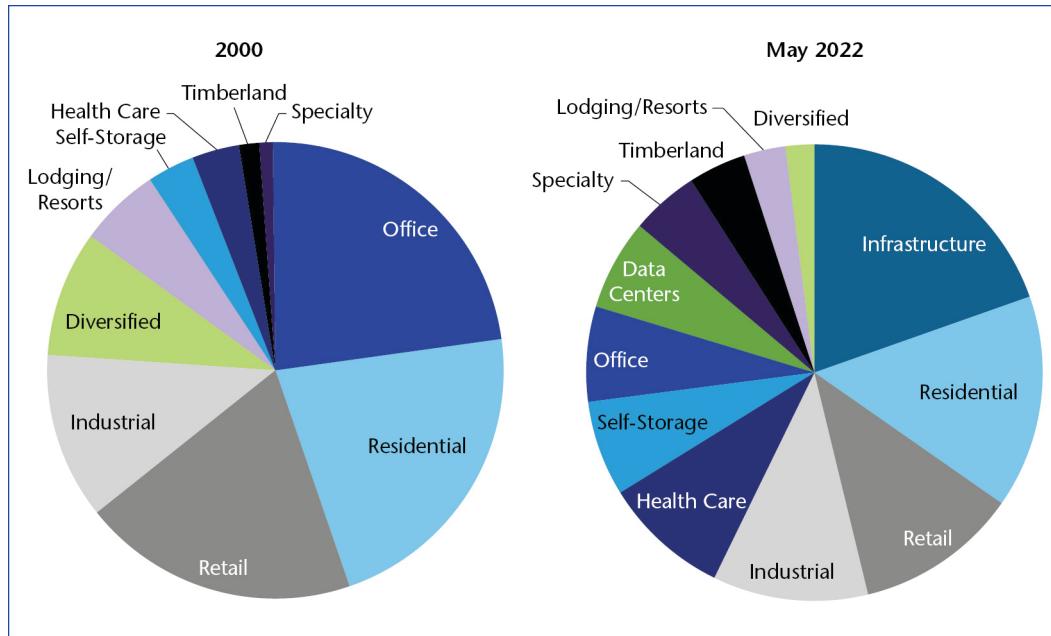


EXHIBIT 26-1 Property Sector Ownership by U.S. REITs, 2000, 2022.

Source: NAREIT, 2022.

categories: office, multi-family residential rental, retail, and industrial. Together, these four categories were good for over 75 percent of the total market capitalization of the U.S. REIT market at that time. This was in line with investor interest at the time: institutional investors' private property portfolios also covered these categories. However, interest has shifted to other "alternative" categories, and this is visible in the REIT market.

By 2022, the U.S. REIT market had undergone some big changes. First, the largest asset types are far less dominant than they were in 2000. Second, the type of assets they hold has changed tremendously. Office and retail have lost their preeminence, and new asset types, such as infrastructure, health care, student housing, single-family rentals, and self-storage, have become more important. Also, the role of diversified REITs has continued to decline. In other parts of the world, similar developments are visible, although less pronounced than in the U.S., and in some countries, diversification across regions and property types is mandatory to maintain the favorable tax status.

This success has been a reason for countries other than the three early adopters to establish tax-transparent real estate ownership structures on the same footing, starting in the 1990s. Gradually, the REIT structure has been adopted far more widely all over the world, with over 40 countries now having some form of tax-transparent real estate ownership structure (EPRA, 2022).

Exhibit 26-2 illustrates how the global REIT market grew over the four decades from the 1980s to the early 2020s. The left y-axis pertains to market capitalization in billions of U.S. dollars, and the right y-axis to the total number of listed REITs worldwide. Starting at approximately \$25 billion in 1983, the total capitalization of the global property share market had reached \$1.9 trillion by the end of 2022, about half of which was in the Americas, with Europe and Asia-Pacific splitting most of the rest. Africa accounts for just under 1 percent of the total market capitalization in 2022. The market development has not followed a smooth trend, but instead has been growing in fits and starts, showing strong growth in boom periods, stability in other times, and a virtual implosion between 2007 and early 2009, when the global market lost more than 60 percent of its value. Although this decline was far more pronounced in North America than in Europe, this was a global phenomenon,

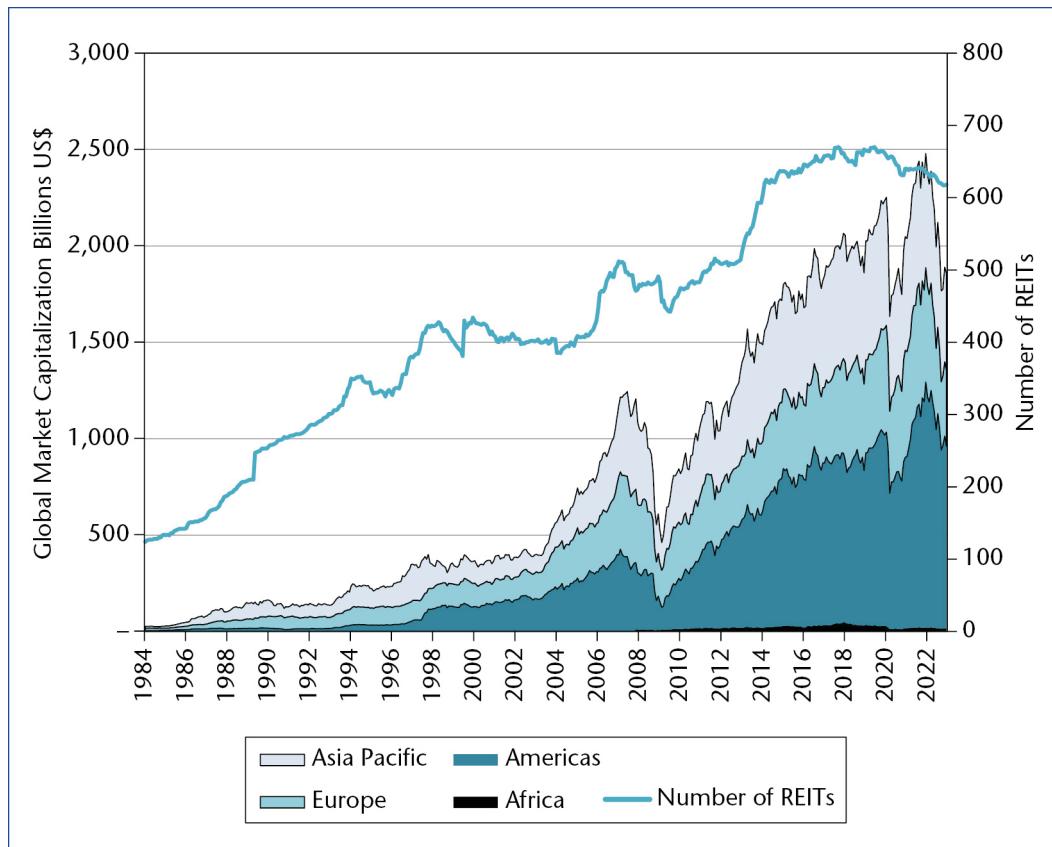


EXHIBIT 26-2 Size of the Global REIT Market.

Source: GPR, 2022.

as was the equally spectacular recovery that followed. Since then, property share markets all over the world prospered until the beginning of the COVID pandemic, which again led to an almost synchronous decline in market values across the three main regions. The recovery again was swift, with the value of the global market reaching nearly \$2.5 trillion at the end of 2021. (As REITs employ debt financing, the implied value of their property asset portfolios is even greater.) Meanwhile, the number of listed REITs has increased from 124 in 1984 to over 600 in 2022, with that number slightly declining in recent years, especially in the Americas and Europe. The size of the average REIT has greatly increased.

Generally, the growth of the global REIT market has four sources, with the first two contributing to the number of REITs and the latter two to their market capitalization. The first source of growth is the IPO market. The 1990s growth in the U.S. REIT market, for example, was mainly fueled by an extended IPO wave: between 1993 and 1997, over 100 U.S. REITs got listed on the stock market. Often, these firms had a history as private companies, and the REIT market provided them with easier access to capital (and in some cases, a tax efficient route to transition from a founding generation of owners). Between 1992 and 2022, there have been 317 REIT IPOs in the U.S. alone, raising a total capital of just over \$78 billion.

Second, when countries introduced REIT regimes, the existing listed property companies tended to convert to the REIT status. This was not always an instant development: the REIT status significantly limits the freedom of management to act as they see fit, so property companies often

retained their regular corporate status. Over time, however, most property companies took the jump and adopted the REIT status to save tax outflow. This development is nicely illustrated in the U.K., where the REIT structure was introduced in 2007, with only 10 listed real estate companies initially making the conversion. Gradually, the others have done so as well.

The remaining two sources of market growth are increased market values of the property portfolios owned by the REITs (and the stock market's valuation thereof), and Secondary Equity Offerings. Especially the latter is a key source of market growth: in the U.S. the total number of REIT secondary offerings (both common stock and preferred stock) between 1992 and 2022 was just under 3,500, cumulatively raising approximately \$0.7 trillion in equity capital.

Over time, these sources led to both a greater number of REITs, and a larger market size for a typical REIT. This converted a market that was hitherto mostly regarded as a vehicle for retail property investors into an institutional investment market, since it could increasingly absorb the magnitudes of large institutions' capital. And this development fed on itself, since the increased availability of capital made it more attractive for property companies to become listed as REITs.

Exhibit 26-2 also clearly shows declines from time to time in the combined market value of the global REIT market. This is the result of declines in REITs' property portfolio value (and the stock market's valuation of it) and REIT delistings. In 2008–2009, REITs fell in value all over the world, but the most notable example of declining values is the U.S. REIT crash around the Global Financial Crisis. In that period U.S. REIT share prices took their most drastic pummeling at least since the early 1970s, with the NAREIT Equity REIT Price Index falling from a peak of 650 in January 2007 to 480 on the eve of the Lehman collapse in September 2008, into a breathtaking spiral down to 180 at the beginning of March 2009: a peak-to-trough fall of 75 percent. Interestingly, REITs were still able to access the equity market with secondary offerings even in 2008, and the 2009 volume of secondary equity capital raised (\$24 billion) surpassed that of 2004, 2005, and 2006, when the REIT market was booming.

This illustrates the mixed blessing of the instant pricing of the stock market: it surely increases the volatility of the REIT stocks, but it also keeps the source of equity capital open. This contrasts with the private property markets, where crises tend to result in strongly reduced access to capital.

26.1.2 OPERATIONAL STRUCTURE AND REGULATORY CONSTRAINTS

The hallmark of the “modern REIT” is that it is a vertically integrated operating company, providing commercial real estate investment goods and services throughout the “real estate value chain.” While not all publicly traded REITs are in fact vertically integrated (or even self-administered), Exhibit 26-3 depicts the general idea and the modern paradigm that does in fact describe many of the largest U.S. equity REITs. However, the extent to which REITs can cover the whole value chain differs across countries, since the right to avoid paying corporate tax is often related to operational limitations, to maintain a level playing field with companies that are required to pay these taxes. Most commonly, property development by REITs is restricted, since pure property development companies (“merchant builders,” which develop property primarily to sell the buildings) must pay corporate tax.

Looking at Exhibit 26-3, the production and delivery of commercial property investment “goods and services” can begin with raw land acquisition and “land banking” (holding land for future development). It continues with the development of structures on the land. Then, the main component of the entire process, which absorbs most of the time and capital opportunity cost, is the holding of stabilized operational property assets as investments. Operational management, including property management, is another important service and can be a major profit center. REITs can offer so-called “downstream” services like tenant services (such as provision of certain utilities and “concierge” type services). This type of vertical integration makes the modern REIT more like a classical “operating company” similar to other typical corporate denizens of the stock market. However, the

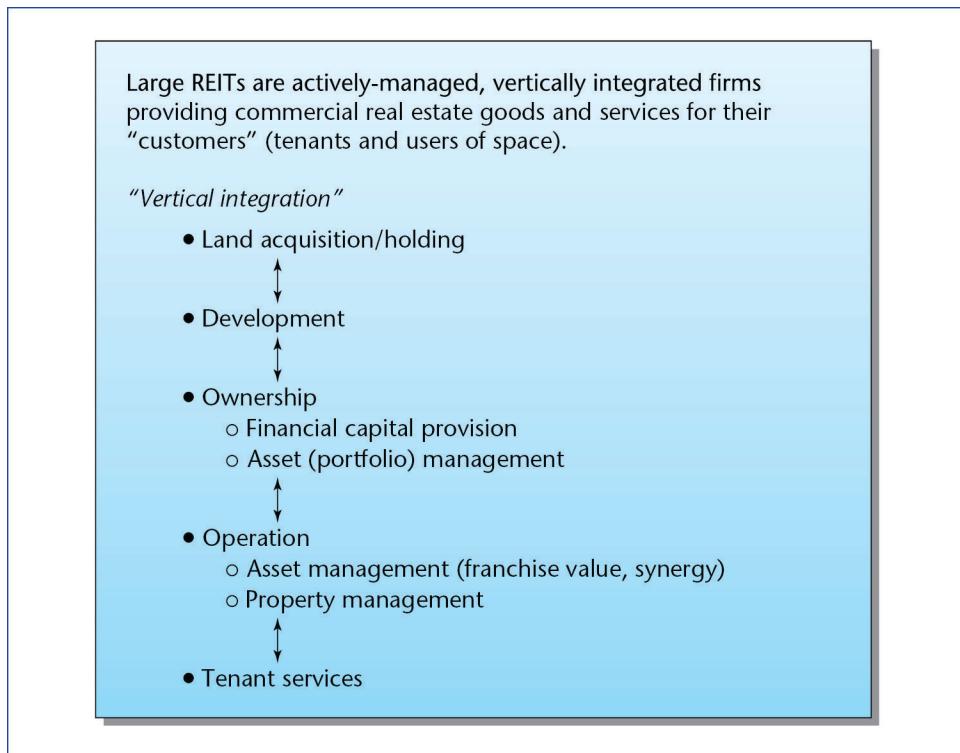


EXHIBIT 26-3 Vertical Integration of REITs.

extent to which this is allowed by regulation differs across countries, and REITs often retain some of their original nature as relatively passive investment funds or vehicles.³

Regulatory Constraints

Perhaps the most notable feature of REITs as compared to most other types of stocks is that REITs are exempt from corporate income tax.⁴ The rationale for this exemption is that REITs are viewed as investment vehicles, similar to mutual funds. Exemption from corporate-level taxation enables REIT shareholders to avoid the double taxation of corporate income that characterizes most stocks. In effect, REITs have tax advantages similar (though not identical) to pass-through entities such as mutual funds and partnerships, while retaining the limited liability afforded by the corporate ownership form. In fact, private real estate investment is usually structured through pass-through entities that also avoid double taxation. So, the REIT tax treatment largely levels the playing field between public listing versus private holding of real estate as far as income taxes are concerned. All over the world, REITs' favorable tax treatment entails regulatory constraints, designed essentially to maintain REITs as a somewhat passive investment vehicle specialized in real estate, not too far removed

³ It is not clear whether REITs create shareholder value by engaging in these non-core activities. Feng and Liu (2023) show that REITs with a higher ratio of rental income to total income achieve higher operational performance and higher market value. However, non-core activities do seem to create value for REIT executives, as these get higher pay when their firms are more involved in other business lines.

⁴ This is commonly stated but not technically true. Dividends paid by REITs are deductible for tax purposes, and REITs must pay out much of their earnings in the form of dividends to shareholders. For most REITs, this implies a corporate tax bill near zero much of the time. On the other hand, REIT dividends received by REIT investors are generally taxable at ordinary income rates (with some exceptions).

from the original “mutual fund for real estate” idea. In the U.S., a property company that wants to qualify as a REIT must meet, on an ongoing basis, the following requirements.⁵

1. **Distribution Test:** At least 90 percent of a REIT’s annual taxable net income must be distributed to shareholders as dividends each year.
2. **Asset Test:** Seventy-five percent or more of a REIT’s total assets must be real estate, mortgages, cash, or federal government securities. Since 2001, REITs have been allowed to form and own a **taxable REIT subsidiary** that allows them to engage in activities and/or services to tenants that were previously not permitted by the IRS under REIT rules. This is limited to 20 percent of their asset value.
3. **Income Test:** REITs must derive at least 75 percent of their income from primarily passive sources like rents and mortgage interest (including partnerships and shares in other REITs), as distinct from short-term trading or sale of property assets. They cannot use their tax status to shield non-real-estate income from corporate taxation. A REIT is subject to a tax of 100 percent on net income from “prohibited transactions,” such as the sale or other disposition of property held primarily for sale in the ordinary course of its trade or business, unless it has held the property for at least four years and the aggregate adjusted basis of the property sold does not exceed 10 percent of the aggregate basis of all assets of the REIT.⁶
4. **Ownership Test:** A REIT cannot be a closely held corporation, in the sense that no five or fewer individuals (and certain trusts) may own more than 50 percent of the REIT’s stock, and there must be at least 100 different shareholders. This is known as the **five-or-fewer rule**. With the **look-through provision** enacted in 1993, pension funds are considered for the purpose of this rule to represent as many owners as there are members of the pension plan. Thus, institutional investors are not limited by the five-or-fewer rule.

Although other countries have adopted regulations that are often similar to those in the U.S., there are key differences across countries. To illustrate that, Exhibit 26-4 provides an overview of tax pass-through property vehicles in a number of key countries: property investment entities that do not pay tax at the corporate level, like U.S. REITs.⁷ The (nearly) common denominator in international REIT regulation is the distribution rule. Although the exact percentage differs across countries, it is nearly always the case that most of the annual earnings need to be paid out as dividends. Other common requirements are similar to the U.S. Asset and Income Tests: most of the assets must be real estate, and most of the revenue should be rental income. Capital structure requirements are also very often present, mostly in terms of leverage ceilings. However, it is also interesting to note some international differences: in some countries, REITs are expressly allowed to engage in residential real estate, while that is limited or prohibited in others. Many countries have specific diversification rules, across properties, regions, or property types, while other countries allow single-asset REITs. Most countries expressly allow cross-border investment by REITs, but some prohibit it. Sometimes, REITs are only allowed to invest in urban real estate, while no such restrictions exist in other countries. In other words, there are large international commonalities, but the devil is in the details.

⁵ The specific rules governing REITs are changed from time to time by the IRS and/or Congress. The National Association of REITs (NAREIT) is an excellent source for updated information.

⁶ This requirement forces REITs to have a longer-term focus when making property investment decisions and implies an average holding period of 10 years with a minimum hold of at least four years. There is some evidence that REITs tend to manage properties for income rather than for price appreciation at the property level, as compared to some other types of investors.

⁷ The European Public Real Estate Association (www.epra.com) provides excellent and up-to-date overviews of listed tax-exempt property vehicles all over the world. The information provided here is from EPRA (2022).

EXHIBIT 26-4

Tax-Transparent Property Structures and Characteristics

Country	Structure	Inception	Debt ceiling	Minimum payout	#REITs
<i>Americas</i>					
Canada	MFT	1994	no restrictions	100% net income	41
USA	US-REIT	1960	no restrictions	90% net income	183
<i>Europe</i>					
Belgium	BE-REIT	2014	65% asset value	80% net profit	17
France	SIIC	2003	thin capitalization rules	95% tax-exempt profits; 70% capt gains	28
Germany	G-REIT	2007	66.25% asset value	90% net income; 50% capt gains	6
Netherlands	FBI	1969	60% real estate / 20% securities	100% net income	5
Spain	SOCIMI	2009	no restrictions	80% rental income; 50% capt gains	74
UK	UK-REIT	2007	interest cover test	90% rental profits	53
<i>Asia/Pacific</i>					
Australia	UNIT TRUST	1985	no restrictions	100% net income	41
Hong Kong*	HK REIT	2003	50% asset value	90% net income	9
Japan	J-REIT	2000	no restrictions	90% distributable profits	63
Singapore	S-REIT	1999	45% asset value	90% net income	39
South Korea	REIC	2001	66% asset value	90% net income	21

Note: * China introduced the C-REIT in 2020, but as of 2022 investments are limited to infrastructure.

Source: EPRA, Global REIT Survey 2022.

26.2 REIT ANALYSIS AND VALUATION⁸

A fundamental and important issue in analyzing REITs is valuation: How much is a given REIT worth? The answer to the question of whether a REIT stock is fairly priced, is obviously important to potential investors in the REIT. But this is also important for REIT managers in addressing all sorts of issues about the strategy and operations of the REIT, because they generally want to maximize shareholder value. This is particularly relevant when making an asset acquisition in the private property market: How will a given acquisition at a given price affect the value of the REIT's stock?

REITs are unique among companies in the stock market in that there are two fundamentally different ways to think about the value of the company. As with most stocks, one can view a REIT essentially as a stream of future cash flows generated by the profitable operation of the company. But one can also think of the firm as a collection of real estate assets that are themselves directly valued in the private property market. Real estate is unique in that its physical capital (property assets) is traded directly or indirectly in two separate, “parallel” asset markets, the private property market and the stock market for REIT equity shares.

Both the stock market and the private property market are long-established and well-functioning arenas for trading capital assets. As we previously noted, capital assets are, essentially, long-lived

⁸ The valuation concepts and calculations introduced in this section are illustrated and integrated within a simple yet comprehensive model, “ABC REIT,” contained in an Excel file in the supplementary materials accompanying the book.

streams of cash flows. Both the stock market and the private property market function as effective information aggregators and equilibrium seekers, although they operate very differently and cater to different investor clienteles. They both provide important information about the value of the real assets they trade.

It is probably correct to say that shares of REITs listed on public stock exchanges are valued essentially the same way other stocks are valued. The U.S. stock market is one of the most efficient branches of the capital market in the world, and so it would be surprising indeed to find REITs segmented off in their own corner of a market that is otherwise highly integrated. For REITs trading in stock markets outside the U.S., less deep and liquid though they are, the same can be said. Within an integrated market, the “Law of One Price” must hold: what is essentially the same thing cannot trade at different prices at the same time. And the stock market generally values stocks by viewing companies as streams of future cash flows. Thus, we will begin by taking this perspective on REITs, and then later turn to the collection-of-assets perspective.

26.2.1 INTRODUCTION TO REIT EARNINGS MEASURES

To analyze REITs as streams of cash flows, we need to begin by giving some background about how earnings and cash flows are measured and reported in the REIT world. While REIT stocks are valued in much the same way as traditional publicly traded operating companies, REIT **GAAP Net Income** (or earnings) is widely viewed as not being directly comparable to the earnings reported for most other types of publicly traded operating companies. Depreciation expenses under GAAP hide considerable amounts of REIT cash flow from being reflected in taxable income. Furthermore, real property depreciation is often not matched by an actual loss in nominal value of the property over time, as inflation may match or even exceed the rate of real depreciation in the property. As a result, the U.S. REIT industry has adopted a special measure of earnings, known as **FFO (funds from operations)**, as a supplement to the traditional GAAP Net Income metric. Alas, FFO does not have a very precise and widely applied definition across all REITs, so it can involve some “apples and oranges.” Nevertheless, to gain a basic understanding of REIT earnings and cash flow metrics and how they relate to traditional private property income metrics, let’s walk through the example in Exhibit 26-5. This is based on the U.S. situation and terminology, but the main takeaways apply to most REIT-like structures internationally.

Suppose we start out with a set of commercial property assets that can produce a potential gross revenue of \$180 million per year if they were fully leased. This is typically referred to in the property world as PGI (potential gross income). Conceptually, PGI exists whether the properties are held in a REIT or not. Let’s see how this set of properties’ income might typically be measured and reported in the private property market world and in the REIT world, assuming the (same) properties were held in a REIT.

First, we subtract the effect of vacancy and operating expenses to arrive at the property’s Net Operating Income (NOI). Supposing typical values as in the Exhibit 26-5, this gets us down to a \$100 million annual operating income bottom line for our portfolio of properties. Most private owners and funds would directly examine and report this property-level bottom-line metric, and indeed many REITs also do report their property-level NOI.

Next, in the case of the REIT we subtract firm-level overhead or what is called “general and administrative” expenses. Assuming “G&A” of \$3 million, this brings us to the REIT’s **EBITDA (earnings before interest, taxes, depreciation, and amortization)** of \$97 million. There is no equivalent deduction or bottom line for direct property income measurement. Of course, in practice an investment portfolio of properties generating \$100 million in NOI would no doubt require some administrative and management functions above the property level, so this is where we begin to look at “apples versus oranges” in the typical private direct property income measures as compared

EXHIBIT 26-5

Widely Used Income Measures in Direct Property versus REITs

Direct Property	REIT	
PGI	180	PGI
Vacancy	-9	Vacancy
Operating Expenses	-71	Operating Expenses
NOI	100	NOI (property level)
		General and Administrative Expenses
		-3
		EBITDA
Interest Expense	-40	Interest Expense
		Funds from Operations, FFO
Depreciation Expense	-20	Depreciation Expense
GAAP Net Income	40	GAAP Net Income
		(Dividends $\geq 0.9 \times$ Net Income = 33.3)
Add back Depreciation	+20	Add back Depreciation
Capital Expenditure	-15	Capital Expenditure
EBTCF	45 ^a	Adjusted Funds from Operations, AFFO
		Funds / Cash Available for Distribution, FAD/CAD
		(Available for plowback: 42 - 33.3 = 8.7)
		Div/FFO as low as 33.3 \div 57 = 58%, Div/FAD = 33.3 \div 42 = 79%

^a Direct property EBTCF is not directly available to a passive investor, whereas REIT AFFO is.

to REIT measures. But as a rough correspondence, considering that G&A is usually rather small, EBITDA in REITs is roughly comparable to NOI in the property world.

Next, we take out interest expenses owed on any debt used to finance the properties. At the direct property level, any such debt would have to be property-level mortgages. But in the REIT world, the relevant debt could be either mortgages or firm-level (“entity-level”) debt such as unsecured bonds or bank loans. Keeping the comparison as fair as possible, we will assume that the interest expense would be the same whether the properties are held in the private property market or in a REIT, so we subtract the same (example) \$40 million from each side in Exhibit 26-5. This subtraction does not produce any particularly widely noted metric in the direct private property world, but it gets us to perhaps the most famous and widely quoted measure of REIT earnings, the aforementioned “funds from operations,” or FFO, which we see is \$57 million in our example.

Now let us suppose that the set of properties underlying the exhibit can charge \$20 million of annual depreciation expense.⁹ Subtracting this, we arrive at the major “official” net income bottom line, which is the so-called “GAAP Net Income.” This is the major measure of profit or earnings for most corporations and is essentially the basis for corporate income taxation. In our example it equals \$37 million for the REIT and \$40 million for the set of private properties computed directly at the property level.

The REIT must pay out at least 90 percent of this GAAP Net Income as cash dividends to its stockholders. Thus, the REIT must pay dividends of at least $0.9 \times 37 = \$33.3$ million. This means that the REITs “retained earnings” in the traditional measure can be no more than \$3.7 million from this

⁹ As a simplified example and with currently prevailing U.S. rules, suppose all the properties are nonresidential commercial buildings (39-year depreciable life) purchased for an historical cost of \$975 million of which 20 percent was land value (hence, depreciable cost basis of $0.80 \times 975 = \$780$ million, which dividing by 39 gives \$20 million/year depreciation). Yet these properties might be worth, say, \$1,250 million (8 percent cap rate applied to \$100 million NOI), and if their market value is appreciating at 1.6 percent per year (which could be quite plausible if inflation is around 3 percent), then none of the \$20 million depreciation expense is actually tracking any loss in value (in nominal terms).

set of properties, which generated \$100 million NOI. But this is not the whole story. So far, we have been constructing accrual-accounting-based earnings measures. Now, we will adjust these measures to get to the actual cash flow, which is more important in practice.

First, we add back the \$20 million depreciation expense since this is not actually a cash outflow. On the other hand, we must now subtract out the capital improvement expenditures made on our set of properties, as these are a cash outflow. Such expenditures are a routine and necessary part of property operations and typically amount to 10 percent to 20 percent of the NOI over the long run. In REITs, as in privately held properties, such routine capital expenditures (aka “capex”) are generally financed internally from the cash flow generated by the property itself. Thus, we subtract out an illustrative \$15 million from both sides of Exhibit 26-5. This brings us down to the essential cash flow bottom line on both sides, what we referred to as EBTCF (equity-before-tax cash flow) at the property level in Chapter 12 and what is often called **AFFO (adjusted funds from operations)** in the REIT world.

AFFO is similar to, and sometimes the same thing as, another measure referred to as **FAD** or **CAD (funds, or cash, available for distribution)**. This refers to the fact that this is in principle free cash flow that is available for the REIT to distribute out to its shareholders if the REIT management chooses to do so, and without jeopardizing the long-run health of the REIT’s existing properties as we have already subtracted out the (presumably adequate) capital improvement expenditures. The major difference between AFFO and FAD or CAD is that the latter two may also subtract nonrecurring or extraordinary capital expenditures, such as for major development projects.

Thus, EBTCF and AFFO are roughly comparable bottom line cash flow measures between the direct property investment and the REIT industries. In our illustrative example in Exhibit 26-5, the property-level EBTCF is \$45 million, and the corresponding REIT AFFO is \$42 million. However, as indicated by the asterisk (*) in the Exhibit, we must not forget that the REIT side is already net of the entity-level overhead and administrative costs of managing the portfolio of properties, whereas we have nowhere taken out such “supra-property-level” costs on the private property side. And note that REIT FFO has no exact analog in the direct property metrics, even though FFO is the most widely cited REIT earnings measure and people often loosely or carelessly try to equate it to the most widely cited earnings metric in the private property world which is NOI. (In fact, the REITs’ EBITDA is more comparable to the private property world’s NOI.) Furthermore, in common parlance, when people speak of “FFO” or “earnings” for REITs, it’s not always clear exactly what they mean (or *should* mean), and they may in fact be referring to something closer to what we have defined in Exhibit 26-5 as AFFO.

There are other differences that we should mention even at this broad-brush level. For example, at the level of generality of Exhibit 26-5, FFO appears very clearly and precisely defined. But the devil is in the details, and there is considerable variation in the real world. Unlike GAAP Net Income (or just “net income” for short), FFO (not to mention AFFO) is not officially and precisely defined. For example, in general, “extraordinary items” of net income, such as income from the sale of properties, are not counted in FFO (in either revenue or expenses). But some REITs that engage more routinely in some sales of their assets (within the constraints of the REIT taxation rules) do include such extraordinary income in their reported FFO.

Some important differences between REIT earnings metrics and traditional property-level metrics result simply from REITs’ requirement to apply GAAP accounting, which is accrual-based, whereas tradition in the private property market has always been more cash flow focused. A notable such consideration affects what the property world calls the “effective gross income” (EGI), the gross revenue generated by the property. In the private property world this is measured as the cash flow from the current leases in place. But in the GAAP accrual accounting world long-term leases must be “straight-lined,” that is, their rents are reported at their average level across all years in the lease contract. For example, a 3-year lease with contracted rents of \$100, \$105, and \$110 in the three

consecutive years would be reported under GAAP rules at \$105 each year. As leases often contain “step-ups,” to increase their rents in later years, this tends to result in GAAP Net Income overstating the cash flow in early years and understating it in later years, for any given lease (though this may average out across many leases signed at different times for different terms). As usually constructed, REIT, EBITDA and FFO will reflect the GAAP-based straight-lining, while property-level NOI and EBTGF will not. However, most constructions of REIT AFFO will adjust to remove the rent straight-lining effect, making AFFO more comparable to EBTGF.

Property-level operating expenses are generally entirely cash outflows in the private market world. But many REITs are in the property management business, at least for their own properties, and some REITs view property management as a profit center. Some of the property management expenses at the property level are profits at the entity level of the REIT. In effect, this reduces the G&A outflow on a net basis.

Returning to Exhibit 26-5, note that an owner of the properties in the private market would in principle be subject to income taxes on the \$40 million of GAAP Net Income, which would presumably be received by private taxable owners through some sort of pass-through entity such as a partnership or fund. The REIT avoids corporate-level income tax on its \$37 million Net Income, but in most countries, its private shareholders must pay income taxes (to the extent the shareholders are taxable) on any dividends they receive from the REIT.

Any portion of the REIT’s cash flow earnings that it does not pay out as dividends is effectively its “retained earnings” on a cash flow basis (or more exactly its “retained cash flow”). For example, if the REIT pays out just the minimum required \$33.3 million in dividends, then it will have retained \$8.7 million for reinvestment in the firm, presumably largely to buy or substantially improve and redevelop its current property holdings (as we’ve already subtracted out \$15 million in routine capital improvement expenditures). In that case, while the REIT will have satisfied its tax obligation to pay out 90 percent of its taxable earnings, it will have paid out only $33.3/42 = 79\%$ of its free cash flow (AFFO), and only $33.3/57 = 58\%$ of its FFO (a commonly cited ratio).

In summary, understanding REIT cash flow and making “apples-to-apples” comparisons to property-level or private-market-based cash flow measures is a nontrivial exercise at the very detailed level. But at the broad-brush level the picture presented here is accurate and sufficient for a useful understanding of REIT cash flows.

26.2.2 VALUING REITs AS A STREAM OF CASH FLOWS: THE GORDON GROWTH MODEL

Let us now turn to the valuation of a REIT as a stream of future cash flows. As noted, this is the traditional way that the stock market values most corporations. It is particularly congruent with the “modern REIT” model described at the outset of this chapter as a vertically integrated operating firm.

Viewed as a stream of future cash flows, the value of a REIT’s equity derives fundamentally from an infinite-horizon discounted cash flow valuation similar to property-level valuation. Now, however, the net cash flows that matter are the dividends paid out by the REIT to its stockholders, instead of just the property-level net cash flow generated by the property assets. Of course, the REIT’s net cash flow for its stockholders must ultimately derive from its property holdings. Letting E represent the current (time zero) value of the firm’s equity (equal to the price per share times the number of shares outstanding), we have

$$E_0 = \frac{DIV_1}{(1+r)} + \frac{DIV_2}{(1+r)^2} + \frac{DIV_3}{(1+r)^3} + \dots = \sum_{t=1}^{\infty} \frac{DIV_t}{(1+r)^t} \quad (1)$$

where DIV_t refers to the annual dividends expected to be distributed by the REIT in year t and r refers to the stock market's required long-run total return expectation for investments in the REIT's shares.¹⁰ The “ Σ ” symbol means summation of the discounted dividends.

While formula (1) is the fundamental stock valuation model, a famous shortcut or simplification of (1), known as the Gordon growth model, is more widely used.¹¹ It follows directly from formula (1) and the constant-growth perpetuity formula introduced in Chapters 2 and 4. The Gordon version of the generalized dividend discount model is obtained by assuming that dividends are expected to grow at a constant rate, g , each year into the indefinite future. Under this assumption, the long valuation formula in (1) collapses into a much simpler formulation of the value of a firm's equity, given by

$$E = \frac{DIV_1}{r - g} \quad (2a)$$

Formula (2a) is obviously a shortcut compared to formula (1) because the analyst does not have to forecast each future year's dividend explicitly. Instead, in the Gordon model, the analyst only has to estimate the market's expected long-term average growth rate in the firm's future dividends, g . Note that $(r - g)$ is essentially a “cap rate” applied to firm-level dividends. The shortcut formula in equation (2a) parallels the cap rate approach to private property valuation, $V = NOI/(cap\ rate)$, that real estate investors use as a shortcut to the more fundamentally sound general discounted cash flow model.

REIT valuation as a stream of cash flows, therefore, typically revolves around estimates of three variables: DIV_t and g have to do with the firm's future cash flows (i.e., AFFO) and dividend distributions, while r has to do with the firm's equity risk as perceived by the stock market, and the stock market's preferences for such risk. As the model assumes a constant growth rate in dividends over the long run, it is only an approximation of reality. To make the model work as well as possible, the variables employed in it should reflect long-run average stabilized and sustainable levels for each variable, the initial annual dividend, the growth rate in those dividends, and the market's expected return on investments in the firm's stock. Temporary aberrations or transient effects in any of these variables can substantially distort the application of the model to any given firm, suggesting at times it may not be advisable to rely solely on the simplified constant growth and constant discount rate model, but augment it with a more nuanced discounted cash flow model or perhaps the REIT as a collection of assets (NAV) approach that we discuss later.

Let us now briefly consider each of the three variables in the Gordon growth model, DIV_t , g , and r , with particular attention to the cash flow growth rate, g , as this is what most sensitively reflects the valuation view of REITs as streams of future cash flows.

¹⁰ This is equivalent to the REIT's average equity cost of capital, expressed in the form of a single (blended) long-run multiperiod required expected return (like a going-in IRR, but at the level of the firm's equity). In previous chapters, we often labeled this long-run required return expectation $E[r_e]$ to emphasize that it is an *ex-ante* expectation of the return to levered equity. Similarly, the DIV_t amounts are expectations. Note that the numerators and denominators on the right-hand side of the valuation equation are measured at the before-tax level from the perspective of the REIT investors. This is convenient, as E refers to the firm's equity market value in the stock market, so DIV_t and r can be related directly to empirically observable public information on the REIT share price and dividends. However, market value is determined more fundamentally as the investment value to the marginal investors. In the case of REITs, these marginal investors will generally be taxed individuals. The more fundamental equity valuation equation would use after-tax dividends and the equity-after-tax opportunity cost of capital for the marginal shareholders of the REIT, though this level is not directly observable empirically.

¹¹ Although the Gordon growth model was not originally invented by Gordon, its modern usage is most often attributed to Gordon and Shapiro (1956).

The initial dividends paid out by the REIT, DIV_1 , is on its face very easy to empirically observe. However, remember that the model must reflect sustainable values. Sustainable dividends must be taken from the net equity cash flow earned by the REIT from its property holdings and sales (i.e., REIT level AFFO plus net sale proceeds, or “FAD”). Thus, stock market analysts and investors must study the nature of the firm’s current property holdings and operations, as well as its debt obligations, as these will largely determine the firm’s ability to pay dividends, at least in terms of the stabilized and sustainable rate for such cash payouts. The cash dividend paid out of REIT earnings is also a function of the REIT’s dividend policy as reflected in the proportion of net earnings it chooses to pay to shareholders.

We can expand the numerator of the model in equation (2a), the expected first-year dividend, by breaking it up into the product of the firm’s first-year earnings and one minus the “**plowback ratio**” p , which is the proportion of earnings retained by the REIT.¹² So $DIV_1 = (1 - p)AFFO_p$, where $(1 - p)$ is the dividend payout ratio, and equation (2a) can be written as

$$E = \frac{DIV_1}{r - g} = \frac{(1 - p)AFFO_p}{r - g} \quad (2b)$$

With this in mind, consider the second variable on the right-hand-side of the Gordon equation, the long-run expected growth rate in the REIT’s dividends, g . Fundamentally, g can only reflect one or more of the following three sources of growth: (i) *growth from assets in place* (existing property cash flow growth as levered); (ii) *growth from investment of retained earnings* (i.e., cash flow from investments made with earnings plowed back into the REIT, possibly combined with additional debt to reflect the REIT’s capital structure policy); and (iii) *future growth opportunities* (reflecting positive-NPV investments, again, possibly as levered).

For now, let’s ignore the third possibility (positive NPV investments). The first growth component, from **assets in place**, also called **same-store growth**, is growth in equity-before-tax cash flow (EBTCF) of the static portfolio of properties consisting of the stabilized operating properties the REIT already owns, as levered by existing debt service obligations.¹³ This, in turn, reflects the growth in the firm’s underlying property (static portfolio) net operating cash flow (PBTCF), as levered by the firm’s existing debt.¹⁴

For example, suppose a REIT holds \$1 billion worth of assets producing \$50 million property-level net cash flow (PBTCF) per year growing at 2 percent per year. Now suppose the REIT’s capital structure is 50 percent debt, such that it consists of \$500 million of equity and \$500 million of debt which we’ll assume is interest-only at a 5 percent interest rate. Then the as-levered same-store REIT returns are (applying the WACC formula (5) from Chapter 12.3):

- Total return: $r_E = r_D + (r_p - r_D)LR = 5\% + (7\% - 5\%)2 = 9\%$, where LR = Leverage Ratio = $1/(1 - LTV)$;
- Growth: $g_E = g_D + (g_p - g_D)LR = 0\% + (2\% - 0\%)2 = 4\%$;
- Same-store cash flow yield rate is: $y_E = r_E - g_E = 9\% - 4\% = 5\%$.

¹² For example, in our previous illustration in Exhibit 26-5, if the REIT paid out exactly 90% of its GAAP Net Income as dividends, \$33.3 million, then its plowback ratio would be: $p = (42 - 33.3)/42 = 21\%$, where \$42 million is its earnings measured as AFFO as in Exhibit 26-5.

¹³ See Chapter 12. This is based even more fundamentally on the property-level operating cash flows described in depth in Part II of this book.

¹⁴ Earlier in the book, we showed that leverage can shift the equity return relatively away from income and towards capital growth (see Section 12.3).

The REIT's same-store growth rate is 4 percent (instead of the 2 percent rate of its property assets), and the REIT's AFFO earnings are \$25 million.

The second growth source relates to the growth in AFFO attributable to EBTCF from new acquisitions or developments financed with retained earnings (as levered), so without issuance of new equity. In general, we would expect such expansion to consist of properties similar in risk/return profile to the REIT's existing asset base, acquired at market value in the private market, or development projects requiring the purchase of land at full market value, in other words, zero NPV deals and projects. Annual expected dividend growth, g , for the REIT in formula (2b) then derives from the previously described growth in same-store property cash flows plus the plowback-based expansion holding capital structure constant, as the REIT reinvests p percent of AFFO each year.

For example, to continue our previous illustration, suppose the REIT pays out \$20 million in dividends and retains \$5 million in equity cash flow (plowback ratio of 20 percent). This will reduce the REIT's dividend yield from 5 percent to 4 percent ($= 20/500$). However, the plowback will enable the REIT's sustainable growth rate, g , to be larger. Suppose the REIT combines the \$5 million equity retained cash with \$5 million in new debt (to keep the 50/50 debt/equity ratio). Also, over the course of the ensuing year the REIT's same-store property assets have presumably (in steady state long-run average expectation) grown by 2 percent from \$1 billion to \$1.02 billion, allowing an additional \$20 million in new debt (again, to keep the D/E ratio at 50/50)

Thus, the REIT invests \$25 million of new debt (\$5 million plus \$20 million as noted) plus the \$5 million retained earnings to acquire \$30 million worth of new property assets (or expansion of existing assets), to give it now (at the end of the year) \$1050 in property assets (the \$1020 that the previous year's assets grew to, plus the \$30 million new assets). As the debt is now \$525 million, the equity is now \$525 million (retaining the capital structure). This is a 5 percent growth rate in the equity (from \$500 to \$525 million). No new shares have been issued, so the growth in value per share is 5 percent, or 100 basis points more than the 4 percent achieved by the same-store assets as levered, due fundamentally to the 20 percent plowback (holding capital structure constant). As noted, the plowback reduces the REIT's dividend yield from 5 percent to 4 percent (by 20 percent), exactly offsetting the growth in share price, leaving the REIT's total return the same as without the plowback, at 9 percent ($r = r_E = 9\%$, which is greater than the underlying properties' 7 percent return, reflecting the REIT's leverage).

We derived the basic Gordon growth model in terms of dividends. But the model can also be applied directly to REIT earnings. Here we have to be careful to employ the relevant growth rate in the denominator, the growth rate in same-store earnings (like the 4 percent rate in the above example) rather than the growth rate in dividends which includes the effect of the plowback (the 5 percent rate in the above example). This means we can write the value of equity as a function of current earnings directly with no explicit consideration of the plowback ratio. Specifically

$$E = \frac{DIV_1}{r - g} = \frac{(1-p)AFFO_1}{r - g} = \frac{AFFO_1}{r - g_E} \quad (2c)$$

with g defined as the long-run growth rate in dividends as above (continuing to assume no NPV > 0 growth opportunities) and noting that this includes the effect of growth due to reinvestment of earnings permitted by the "plowback" of earnings, and g_E defined as the expected growth rate on REIT earnings (i.e., AFFO) from pre-existing or in-place assets only (as levered). Notice that the value of the REIT's shares is independent of the plowback ratio here.¹⁵ In this world, shareholders are indifferent between a passive REIT that owns a static portfolio of properties and pays out all

¹⁵ For equation (2c) to hold, it must be that $(1-p)(r - g_E) = (r - g)$, which solving for g yields $g = g_E + p(r - g_E)$. Intuitively, g , which includes the plowback reinvestment effect, is comprised of two parts: growth from existing assets and growth from reinvestment, with the latter given by the levered income yield (y_E) on property equity, $(r - g_E) = y_E$ times the plowback rate: $g = g_E + py_E$.

property EBTCF as dividends, and an otherwise similar REIT that retains some part of AFFO above the 90 percent of Net Income that must be paid out and expands through acquisitions of properties similar to the ones it currently owns. Of course, the crucial assumption here is the absence of $NPV > 0$ projects. If we would allow for such projects, and if the REIT's managers would be better at identifying them than the REIT's shareholders, earnings retention would add value to shareholders.

In some respects, it is more realistic to apply the Gordon growth model to AFFO rather than to dividends. Most publicly traded companies that pay dividends do not maintain a fixed dividend payout ratio as we have assumed here but instead follow a stable dividend policy, maintaining a roughly constant dollar value of dividend payout that may be increased from time to time to a new higher level but only once the firm feels that it can permanently maintain the higher dividend. Over time, the cash dividends paid have a stair-like pattern for such firms, assuming significant earnings growth over time.¹⁶ Dividend policy may change, but by itself and within reason, that will not change the value of the firm. Therefore, focusing the analysis on earnings and same-store growth can be more accurate and can also provide a sort of "reality check," as ultimately the REIT's value can derive only from the current assets in place plus positive NPV growth opportunities.

To summarize up to here using our previous simple numerical example, the REIT's dividend growth rate, g , was built up starting from the inherent growth in nominal value of the REIT's same-store assets (which we assumed as 2 percent) using leverage (which bolstered the growth in same-store earnings to the 4 percent rate that we labeled g_E). The REIT's dividend growth rate was further augmented to 5 percent by the plowback of retained earnings holding capital structure constant. So far, none of this implied any nonzero NPV investment or any expected returns differing from the market's equilibrium return expectations in the private property market (7 percent unlevered, 9 percent levered 50/50). Furthermore, the REIT's earnings/price ratio (inverse of its "P/E," which in our terminology is E/AFFO) is $25/500 = 5$ percent ($P/E = 20$). In this case the earnings yield of 5 percent is the same as the REIT's property-level cash yield of $50/1000$ (PBTCF/property value), because the cash yield on the interest-only debt is also 5 percent.¹⁷

The REIT can also grow by issuing new shares in a secondary public offering, an SEO. Issuing new shares will increase the "float" (the number of shares outstanding) but does not in itself change the price per share, assuming the REIT can only invest at $NPV=0$. The REIT can also expand in scale by issuing more debt. But increasing the debt/equity ratio will increase the risk in the REIT's equity, and still will not increase the value of the REIT's shares if the capital is all invested at $NPV=0$. Thus, tapping external capital in itself simply to grow the REIT's scale will by itself not increase the REIT's price/earnings ratio.¹⁸

This brings us to the third component that can affect the stock market's perception of the firm's ability to grow dividends in the long run. This is the ability of the REIT's management to obtain and effectively implement **growth opportunities** by making *positive NPV* investments. Typical sources of such growth opportunities for REITs include:

¹⁶ Publicly traded firms generally do not want to cut dividends, as this tends to be viewed by stock investors as a negative signal about the firms' earning prospects. Consistent with this notion, Bradley, Capozza, and Seguin (1998) show that REITs with relatively higher volatility in earnings choose lower dividend payout ratios to reduce the possibility of forced dividend cuts. In line with that, Boudry (2011) shows that REITs use the discretionary component of their dividend to engage in dividend smoothing. On the other hand, REITs that cut dividends during the financial crisis of 2008-2009 experienced positive cumulative abnormal returns after the announcement (Case et al., 2012).

¹⁷ Recall the leverage ratio version of the WACC: $y_E = y_D + (y_p - y_D)LR$, where the leverage ratio $LR = 1/(1-LTV)$.

¹⁸ A possible exception to this "efficient market" result could occur if augmenting the REIT's scale somehow in itself adds value. This possibility will be discussed in Section 26.3. Other nuances on the essential efficient-market paradigm described here could result from direct effects of capital structure or dividend policy, such as signaling effects, or if the issuance of new debt creates too much leverage, adding to the potential cost of financial distress or jeopardizing the REIT's management flexibility.

- Developable land *already owned*
- Entrepreneurial abilities (in development, negotiation, or possibly in other activities)
- Firm-level opportunities (such as economies of scale, franchise value, economies of scope);
- Differential micro-level asset pricing between the private property market and the stock market (often referred to as cross-market “arbitrage,” to be discussed further in Section 26.2.4).

If we assume that the market now expects management of our previously rather passive REIT to be able to consistently create value by these types of opportunities, then the value of the REIT’s shares will exceed that given by the Gordon model formulas (2b, 2c), by the sum of the expected NPVs from the firm’s growth opportunities. In this case,

$$\begin{aligned} E^* &= \frac{(1-p)AFFO_1}{r-g} + NPV \text{ (growth opportunities)} \\ &= \frac{AFFO_1}{r-g_E} + NPV \text{ (growth opportunities)} \end{aligned} \quad (3)$$

where $E^* > E$, with E our previous valuation ignoring positive-NPV growth opportunities. We can rewrite equation (3) in the Gordon model form by directly incorporating the expected positive NPV growth opportunities into the presumed dividend growth rate,

$$E^* = \frac{DIV_1}{r-g} = \frac{(1-p)AFFO_1}{r-g^*} \quad (2d)$$

where g^* reflects dividend growth from all three components we have discussed: same-store growth in EBTCF, growth in EBTCF from new ($NPV = 0$) acquisitions financed with retained earnings, plus growth from positive NPV investments. Exhibit 26-6 summarizes the Gordon Growth Model under the different cash flow growth scenarios we have considered.

Much of the fundamental volatility in a REIT’s share price probably derives from fluctuations in the market’s assessment of these three components of the long-term growth g^* . Holding the other two variables in the Gordon Growth Model formula constant, even relatively small variations in the market’s perception of the REIT’s long-run growth rate can have a large impact on REIT value. For example, if g^* increases from 2 percent to 3 percent while the REIT’s cost of capital remains at 10 percent, the REIT price as a multiple of its current annual dividend will increase from 12.5 to 14.3, a jump of over 11 percent in share value.¹⁹

Finally, let us turn to the third of the three variables on the right-hand side of the Gordon valuation model of REIT cash flows, the discount rate, r . This variable represents the market’s required expected long-run total return, the REIT’s **average equity cost of capital**. Like the growth rate, the discount rate is also very sensitively affected by the market’s perception of REIT management’s long-run abilities. As always, $r = f + RP$, so the market’s required total return expectation equals the current risk-free interest rate (which is exogenous to the REIT) plus the market’s required risk premium for the given REIT. Thus, the market’s perception of the amount and nature of risk in the firm’s future dividends and share value is crucial in determining r .

Fundamentally, the risk in the firm consists of some combination of the risk in the firm’s existing (levered) assets plus the risk in the firm’s management and non-asset-based operations, including taxable REIT subsidiary activities and risk in any future positive-NPV growth opportunities the

¹⁹ The firm value goes from $E = DIV_1/(0.10 - 0.02) = DIV_1/(0.08) = 12.5 DIV_1$ to $E = DIV_1/(0.10 - 0.03) = DIV_1/(0.07) = 14.3 DIV_1$.

EXHIBIT 26-6

The Gordon Growth Model of REIT Share Prices; A Summary

Assumptions About Growth Environment	Comments
Case 1: No Expansion and No Plowback $DIV_1 = AFFO_1$ $E = \frac{DIV_1}{r - g_E} = \frac{AFFO_1}{r - g_E}$	REIT is a passive, pass-through entity owning a static portfolio of properties. DIV growth, g_E , is growth in EBTCF from existing assets in place; growth in same-store levered property income.
Case 2: Internally Financed Expansion but No Growth Opportunities $E = \frac{DIV_1}{r - g} = \frac{(1 - p)AFFO_1}{r - g} = \frac{AFFO_1}{r - g_E}$ $0 < p < 1$ $g = g_E + p(r - g_E) = g_E + Py_E$	REIT grows by reinvesting $p\%$ of AFFO each year; DIV is less than AFFO. REIT buys properties identical to the ones it currently owns, at market value (i.e., NPV = 0), using retained AFFO and debt, keeping leverage constant) REIT equity value is unchanged from case 1. DIV growth, g , exceeds same store EBTCF (and AFFO) growth, g_E , but the REIT's price/earnings (E/AFFO) ratio is the same as in Case 1.
Case 3: Internally Financed Expansion and Growth Opportunities $E^* = E + NPV(growth\ opportunities)$ $E^* = \frac{(1 - p)AFFO_1}{r - g} + NPV(growth\ opportunities)$ $= \frac{AFFO_1}{r - g_E} + NPV(growth\ opportunities)$ $E^* = \frac{DIV_1}{r - g^*} = \frac{(1 - p)AFFO_1}{r - g}$	REIT is perceived to have the ability to find and execute $NPV > 0$ deals or projects, possibly at times due to differential pricing in public versus private real estate markets. g^* incorporates future increases in AFFO due to positive NPV growth opportunities into the growth rate in the Gordon growth model; it "merges" the impact of growth opportunities into g_E , thus $g^* > g$, and the Case 3 REIT's price/earnings ($E^*/AFFO$) ratio is greater than that of Cases 1 or 2.

REIT might face. In effect, much of the firm's fundamental equity risk may be attributable to the market's perception of the risk in the firm's management ability, including its ability to handle leverage. In principle, equilibrium asset pricing models such as the CAPM can be used to help estimate the market's required expected risk premium, as this is applicable to a given firm or, more realistically, to a class of comparable types of firms.²⁰

To summarize this discussion of the cash flow-based valuation of REITs, it should be clear that this is a challenging exercise. A typical REIT is far more complex and dynamic than an individual property, and so forecasting and evaluating it as a future cash flow stream is much more difficult. Yet REITs are relatively simple and transparent types of companies compared to most corporations that trade in the stock market. REITs at least are limited in the type of business they can do, and even in their dividend policies and ability to trade assets. And the type of business REITs can do is relatively simple and predictable as compared to many other industries in the modern stock market. Real estate is not "rocket science." Stock market analysts have traditionally valued most, if not all stocks using the Gordon Growth Model, including many corporations that are much more complicated, dynamic,

²⁰ REITs tend to provide higher returns than predicted by the traditional, single factor CAPM, as they seem to have a "value" stock component to their returns. Peterson and Hsieh (1997) applied the Fama-French (1992) three-factor model to REITs, augmenting the market factor with size and book-to-market (or value) factors. Subsequently, Derwall et al. (2009) showed that the fourth (momentum) factor is of particular importance for REITs. It is now common practice to apply these three- or four-factor models when analyzing REIT returns.

and unconstrained than REITs. REITs do not present anything particularly different or challenging at the fundamental level for the traditional stock market analyst world. Cash flows are cash flows. However, REITs can also be valued using a very different approach, to which we now turn.

26.2.3 VALUING REITs AS COLLECTIONS OF ASSETS: PARALLEL ASSET MARKETS AND NAV-BASED VALUATION

Analysts and investors do not need to value REITs only from the perspective that they are streams of future cash flows. Since REITs' real estate assets trade directly in a well-functioning asset market (the private property market), we can calibrate the outputs from the cash flow-based analysis by valuing the REIT as the portfolio of property assets that it holds.

This approach is straightforward: estimate the value of all the properties currently held by the REIT, in the way these properties would currently be valued in the private property market. Then adjust for non-asset-based value the REIT might have, as from taxable REIT subsidiary operations such as third-party property management services. (In most REITs, this is usually a minor consideration in the overall firm value.) Then subtract the value of the REIT's current liabilities, including claims of preferred stockholders. This gives you the **net asset value (NAV)** of the REIT. Dividing by the number of shares outstanding gives the REIT's NAV per share.

The main task in the above-described procedure is the estimation of the values of the REIT's property assets. However, most REITs hold many properties, and while listed REITs in many countries provide asset-level financial information, it isn't practical to individually appraise or estimate the value of each property. A type of "mass appraisal" or shortcut estimate of value must be performed. This may involve segregating the REIT's EBITDA (or better yet, its property-level NOI) by market segment (such as the type and location of the properties), and then applying an estimate of typically prevailing cap rates in the private market.

In effect, one is performing a valuation of future expected cash flow streams, only at the property level and as evaluated by the private property market rather than at the REIT level as evaluated by the stock market. The process is inevitably only an approximation. Under GAAP accounting rules, REITs in the U.S. are not required to mark their assets to market.²¹ Thus, U.S. REITs carry their assets on their books at their historical cost (less accumulated depreciation). There may be a tendency to be influenced by the REIT's share price, or by the latest "buzz" and excitement in the private market. No two estimates of a given REITs NAV made by different analysts are ever exactly the same.

The REIT's NAV/share can be directly compared to the REIT's current stock price. Typically, this is done by the ratio of share price divided by NAV, or "P/NAV." Any difference of P/NAV away from 1.0 must then be considered. Is it a legitimate and "correct" difference reflecting genuine aspects of the REIT's value not included in the NAV computation? For example, does a P/NAV > 1.0 reflect the REIT's management's ability to add value? This implies an ability to identify and execute positive-NPV opportunities. What are those opportunities, and are they really valuable enough to justify the premium of P/NAV > 1.0 ? If not, then the REIT may be overvalued in the stock market. On the downside, does a P/NAV < 1.0 justifiably reflect poor management's destruction of value? If not, then the REIT may be undervalued in the stock market. And if the P/NAV < 1.0 is justified, then what can be done to improve the REIT management or improve its situation that is destroying value?

²¹ This is in contrast with most other countries, where international accounting standards rely more on the principle of "fair value" than historical cost. Under fair value accounting, which the United States is gradually moving towards, REITs may engage more in their own (or their accountants') estimates of the market values of their properties.

26.2.4 REIT PUBLIC/PRIVATE ARBITRAGE

REIT valuation as a collection of assets is an important exercise at the specific level of analyzing individual REITs. But it can also be aggregated to an industry-wide level, where it may hold additional interesting implications. In particular, industry aggregate or average P/NAV ratios may suggest pricing “disagreements” between the public stock market and the private property market. Such differences, if they are substantial and persistent, can hold implications for public/private “arbitrage” in one direction or another.

The key point to reiterate here is that the stock market valuation of property and the private property market valuation of property are not always consistent. The stock market sometimes values properties more highly than the property market does, not just for a few specific REITs or a few specific properties, but generally and widely. At other times, the reverse is true. In the long run, the two markets tend to agree. *The ability of capital to flow from one market to the other ensures that such agreement will tend to occur on average over the long run.* But in the short run, often, in historical experience for periods lasting up to several years, the two markets may disagree about value. This disagreement can occur in either direction, but in general, when it occurs, the REIT market tends to be ahead of the private property market: after a period typically ranging from one to three years in historical experience, the property market often follows in the direction in which the REIT market previously led.²²

When the stock and property markets disagree about property value, REITs face either positive or negative NPV opportunities from buying properties at market value in the property market. Indeed, there have been prolonged periods in which REITs could implement positive NPV investment strategies either by buying or by selling in the private property market. When the stock market values property more highly than the private property market does, REITs can grow merely by buying properties, and this can turn most REITs temporarily into growth stocks.

Growth stocks are stocks of companies that tend to produce a relatively large share of their total return from capital growth (through share price appreciation) rather than through income (dividends). This requires that such companies be able to do positive-NPV investments or actions. Therefore, growth stocks tend to have higher price/earnings multiples than **income stocks**. Income stocks are similar to and often the same things as **value stocks**. Normally, most REITs tend to be value stocks or income stocks rather than growth stocks.

When the private property market values properties more highly than the stock market does, REITs can, in principle earn positive NPV profits for their shareholders by selling properties into the property market. In such periods, REITs need to become “shrinking stocks” for the benefit of their shareholders, although they need not necessarily get out of the real estate *operating* business. REITs may go private via “leveraged buyouts” (LBOs) in which their shareholders are bought out using money borrowed by or on behalf of the REIT, and then the REIT is sold into the private market either in whole or piecemeal, for value in excess of the money borrowed to pay off the REIT’s public shareholders. Alternatively, REITs may set up private investment funds and effectively sell some of their properties entirely or partially into these private funds (at private market values), with the REIT still running the fund and managing the properties. Thus, REIT “public/private arbitrage” can be performed in both directions, and this can be a source of positive NPV. However, this is not exactly arbitrage as traditionally defined, as it still involves risk.

²² Furthermore, while the stock market for REITs shares generally seems to have more effective and timely price discovery than the private property market in the aggregate, the opposite may be the case at the disaggregate level across individual REITs. Gentry, Jones and Mayer (2004) applied Green Street Price to NAV ratios to theoretically construct a cash-neutral investment portfolio of REITs, buying REITs whose P/NAV was less than 1.0 using funds from shorting REITs whose P/NAV was greater than 1.0, cashing out and reinvesting monthly. This trading rule appeared to be abnormally profitable even after accounting for likely transaction costs, which suggests that the private property market contains information not yet reflected in individual REIT prices.

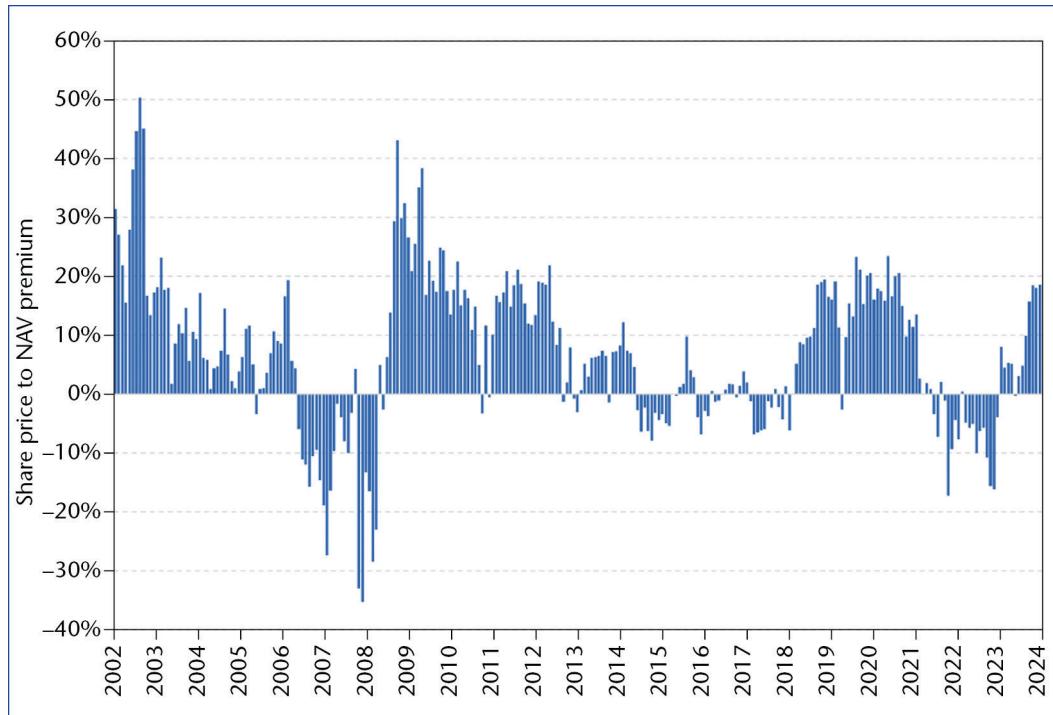


EXHIBIT 26-7 Green Street Average U.S. REIT Share Price to NAV Premium, 2002–2024.

Source: Green Street Advisors, Inc., 2024.

As a result of the shifting relative valuations of property by the stock market and the private property market, REIT investors have to cope with the fact that REITs can change their stripes, so to speak, from being growth stocks with high price/earnings multiples, to being income stocks with low price/earnings multiples, to occasionally being shrinking stocks with very low price/earnings multiples. Predicting these turning points in REIT valuation is difficult, which is a source of risk in REIT investment. However, the private property market is relatively predictable, and this helps to make possible the profitable redeployment of capital between the two types of markets, for some investors.

Exhibit 26-7 displays the average REIT share price premium to NAV reported by Green Street Advisors. It clearly highlights the differential public versus private market pricing for real estate, and it suggests that the premium follows a mean reverting process. That is, there is persistence in either the premium or discount over periods of a few years and then reversal from one to the other.²³

26.3 REIT MANAGEMENT STRATEGY

We already noted the traditional real estate cliché that only three things matter in real estate: location, location, and location. In the 1990s, at the outset of the development of the modern REIT industry characterized by internally managed, vertically integrated corporations, REIT promoters

²³ Kim and Wiley (2018) show that this NAV premium is predictive of U.S. REITs' property investment activity, with increased investment levels in periods of a stock price to NAV premium. Unfortunately, REITs do not seem to exploit the possibilities for value-adding investment to the full extent, as they significantly overpay in their private property transactions relative to other buyers in these periods.

and analysts were fond of paraphrasing this, claiming that three things mattered in the successful REIT: *management, management, and management*. Like most clichés, what makes this one amusing is that it contains an important grain of truth. In this section, we aim to note briefly several REIT strategic management considerations or goals that are often important.

REITs in the “Land of Growth” The stock market tends to like growth and to reward firms that it perceives as able to grow in share price over the long run. Growth firms are awarded with higher price-to-earnings multiples in their share prices. While the reason or rationale is still a bit of a puzzle, the way the stock market has behaved during much of its history, growth stocks have had higher price-to-earnings multiples not just because a larger fraction of their expected total return is in the capital appreciation component, but because these stocks actually tend to provide a lower expected total return. This may be more than just the capitalization into present value of future positive-NPV opportunities. In other words, growth stocks often really appear to be more “expensive” than value or income stocks. In part, this is just a reflection of the Fama-French stock pricing model that we described in Chapter 23, where we saw that stocks with high book-to-market ratios require an extra “risk premium” in their expected returns. The lower book-to-market ratios of growth stocks are the flip side of this.

What does this mean for REITs? Real estate is a “cash cow,” not a “growth rocket.” Whether growth is defined by positive-NPV investments or by simple scale expansion, REITs are at a bit of a disadvantage compared to many other sectors of the stock market. The stock market’s low preference for income or value stocks may be good news for investors in such stocks as it tends to provide them with higher returns. But it is bad news for the companies issuing such stocks because it increases their cost of capital. Keep this challenge in mind as you read about the various strategies described below for how REIT managers try to maximize the share prices of their firms.

Financial Strategy A key component of overall firm-level business strategy is its financial strategy. This is particularly true in the case of a capital-intensive industry such as real estate. Regarding REIT capital structure, it is important to point out some considerations that are particularly noticeable in the financial behavior of REITs. The first point to note is that REITs’ 90 percent earnings distribution requirement may tend to put some REITs under added financial stress. The distribution requirement means that, in order to grow, REITs must go to external sources of capital more frequently.

This situation is probably made more acute by the stock market’s preference for growth. It is therefore natural for REITs to want to try to grow, and this requires infusions of external capital into the firm. Such capital comes generally in the form of either debt or equity. Equity generally comes with less risk to the pre-existing shareholders and fewer strings attached than debt, but it may risk diluting existing shareholders’ value if the REIT’s stock is currently priced low. In addition, REITs do not have the need for corporate income tax shields that most profitable corporations have, which causes those corporations to prefer debt over equity. For these reasons, we might expect that REITs would tend to prefer equity over debt financing.

Yet this natural preference for equity is mitigated to some extent by the growth motivation. Financial leverage can often be employed to convert income-oriented, low-growth underlying assets such as stabilized income properties into higher-growth (lower-yield) equity investments. If a REIT can borrow at interest rates lower than its properties’ total returns, then positive leverage will exist in the REIT’s equity.²⁴ Indeed, if the REIT can borrow at debt constants (including amortization as well as interest) lower than its properties’ current income yields, then it will be able to shift its equity total

²⁴ Note that such an increase in the REIT’s equity total return expectation does not necessarily (or generally) imply an increase in the REIT’s overall firm-level cost of capital, which is measured by its weighted average cost of capital (WACC), including the cost of the debt capital.

return relatively toward the capital gain component without sacrificing current yield. These financial leverage considerations may give debt financing considerable appeal to REITs, even though REITs lack the tax-driven appetite most corporations have for debt financing.²⁵

But if REITs have a desire for financial leverage, this desire tends to be kept in check by the stock market's squeamishness about REIT debt. REITs have gotten into trouble in the past as a result of being too highly leveraged, most famously in the "REIT debacle" of the early 1970s, and more recently again in the financial crisis of 2008. Excessive leverage played a prominent role in the widespread commercial property financial problems at the end of the 1980s and again in the 2000s. As a result of these memories, perhaps, the stock market has the reputation for severely penalizing REITs that find themselves with large amounts of debt relative to the stock market's valuation of their equity. For the most part, therefore, REITs try to retain debt amounts no greater than about 50 percent of their total asset value. This is less debt than is typically employed by taxable privately held real estate firms, but more than most other publicly traded firms.

This relatively conservative debt position, combined with the pressure for growth, makes it easy for REITs to be tempted by a strategy of using low-interest debt to maximize the degree of positive leverage. This typically implies a temptation to borrow using short-term and floating or adjustable-rate debt. Of course, there is no free lunch. This type of debt magnifies the REIT's equity risk and can result in an asset/liability maturity "mismatch" problem. Since this is an obvious danger, REIT analysts tend to be on the lookout for it, and responsible REIT managers resist the temptation. In short, REIT financial strategy is often a balancing act, a tension between opposing objectives and constraints.²⁶

Specialization: Investment Focus and Firm Value A prominent characteristic of REIT management strategy during the 1990s was the tendency of the new REITs to specialize or focus on one type of property, or sometimes two closely related types, such as industrial and office properties. Earlier REITs had often been diversified by property type. The idea was that the REIT's investors would want a diversified portfolio of properties. This may have made some sense when REITs were viewed as passive investment vehicles, mostly catering to small private investors. However, the modern REIT investor community is mostly institutional in nature, preferring to make diversification decisions themselves. Furthermore, when REITs become actively managed entities, it is obvious that management expertise can be more effectively employed when it is specialized by property type. That way, management can harvest returns to scope, for example by catering to the same retail or office tenants all over the country, and by accumulating know-how on optimal building design and location choice in the light of these tenants needs. That would be less obvious with a portfolio consisting of a wide variety of property types.

Also, perhaps more importantly, the stock market can more easily understand and analyze a REIT that is specialized in one of a few somewhat standard space market segments.²⁷ Led by healthcare REITs in the late 1980s, REITs in the 1990s tended to specialize in fields such as apartments, shopping malls,

²⁵ REITs may still face an indirect tax-based incentive to borrow, because of their 90 percent dividend payout requirement. Debt shields REIT taxable income, and therefore reduces the amount of dividend that must be paid out to shareholders, and REIT shareholders dividend income is subject to taxation at the personal level. In addition, there are other, non-tax-related arguments both for and against debt financing, as discussed before.

²⁶ In general, the academic literature suggests REIT capital structure is heterogeneous and not well explained theoretically, with the lack of strong tax or signaling motivations giving REITs more freedom to choose alternative structures. Boudry et al. (2010) and Ertugrul and Giambona (2011) provide evidence and discussion. Dogan et al. (2019) give international evidence, and show that REIT regulations, such as outright leverage restrictions, play a key role here. Besides that, they find that higher dividend payout requirements are associated with higher debt ratios.

²⁷ Capozza and Seguin (1999) provide evidence that the primary value of specialization is not added value in the firm's operating profitability but rather in the stock market valuation of the firm. However, the academic literature does not provide a clear consensus when it comes to the (dis)advantages of REIT specialization. However, Ro and Ziobrowski (2011) counter the notion of specialization value, as do Gyourko and Nelling (1996) and Ambrose et al. (2000), and Demirci et al. (2018) show that property type diversification is associated with lower cost of debt capital for REITs.

smaller retail centers, hotels, and office and industrial properties. This trend seems permanent as it still continues, not just in the U.S., but globally. Over the last two decades, smaller niche products have also been developed, such as self-storage, manufactured housing, student housing, and life science real estate. And the REIT structure has also expanded into infrastructure, including cell-towers, for example.

Vertical Integration Vertical integration refers to the concept of a single firm controlling several linked stages in the production process, for example, from iron ore to steel to cars in the classical example of the Ford Motor Company early in the twentieth century. The old REITs before the 1990s were confined largely to one phase of the process of the production of commercial property goods and services, namely, the role of capital provision (investment) and ownership of existing income-producing properties, which sometimes included some aspects of the role of portfolio or asset management. Vertical integration involves going “upstream” in the production process to the construction and development of new buildings, and even the acquisition of land sites for future construction. It also involves going “downstream” in the production and delivery process to include property management, leasing, and other related services.

A key advantage of vertical integration is that it gives a REIT flexibility to survive, and even profit from, the changes in the relative valuations between the stock and property markets. For example, vertical integration allows REITs to profit by selling properties when their NAV exceeds their stock market valuations, while retaining operating scale and geographical scope by continuing to control the operational management of the properties they sell. Most modern REITs exhibit at least some degree of vertical integration, and this has become a hallmark of modern REIT management strategy. However, the degree to which this is allowed by the tax authorities differs very strongly internationally. Often, REITs can only maintain their zero corporate tax status when they stick to the management of a standing investment portfolio.

Economies of Scale: Is Bigger Better? Scale economies are present when average production costs decline with increasing production. Many capital-intensive industries exhibit scale economies, and one may argue that REITs are subject to this same phenomenon. In the case of REITs, average cost might be measured as the total cost (including operating costs and the required return on capital) per square foot of space offered to tenants (or per dollar of rental revenue). If economies of scale exist in REITs, then REIT overall expenses should increase at a decreasing rate as size increases. Supporters of the bigger-is-better hypothesis suggest that scale economies exist in REIT-level expenses and also in the cost of capital. The implication is that the big REITs will tend to get bigger until the scale economies are exhausted. This suggests that the industry will tend to consolidate, as happened in other capital-intensive industries in the twentieth century. If this is true, then REITs can “grow by growing,” in a sense. In other words, REITs may face positive NPV-opportunities from routine expansion (i.e., even when there is no positive NPV from the investment at the *micro-level*) simply because the increased scale will allow their average total costs to decline. This would obviously have important strategic implications, not only for individual REITs but also for investors interested in the industry as a whole. It also provides an argument for raising external capital.

But is it true? Do REITs really face economies of scale, and if they do, are these economically significant, and until what size of firm do they play out? Is there an optimal scale for REITs? These questions have yet to be answered definitively, although some serious academic research has tried to. Empirical studies consistently find evidence of scale economies in REIT administrative and management expenses, as well as in revenue generation, although the advantages appear to be rather small.²⁸ One would also expect scale economies in capital costs, since larger-capitalization stocks

²⁸ For operational efficiency, see Capozza and Seguin (1998). Ambrose et al. (2005) report evidence of economies of scale particularly with respect to REIT overhead expenses; increasing REIT size lowers average expenses and increases profit margins. They also find that larger REITs have lower costs of capital. Ambrose et al. (2019) find evidence of scale

are more liquid and more likely to be part of the leading equity indices, which would create more institutional investor appetite. However, evidence for that notion is weak.²⁹

Perhaps the most compelling evidence in support of the “bigger is better” hypothesis comes from the tendency toward consolidation in the REIT industry, now evident through a couple of REIT cycles. REITs, on average, have become much larger in terms of market capitalization. Looking back at Exhibit 26-3, the industry has been consolidating since its initial 1990s IPO boom. This trend is visible in many countries outside the U.S. as well. Internationally, the average publicly traded REIT grew from a market capitalization of \$212 million in 1984 to \$3,000 million at the end of 2022 (GPR, 2022). This is almost a fourteen-fold increase in nominal market value. Many REITs have moved out of the small cap realm and are considered mid- and even large-capitalization stocks, with a number being part of banner stock indices such as the U.S. S&P 500 Index or the U.K. FTSE 100 Index.

Branding Another strategy that some REITs began to pursue seriously in the 1990s was to add value by building brand name recognition and reputation among their ultimate customers, the users of their spaces. This strategy has worked successfully for some firms in other branches of the service sector, such as retail and restaurant chains. For real estate, it seems most appropriate for products catering to consumers, such as retail, self-storage, and (student, senior) housing. A notable example of building brand image was that of Simon, a shopping mall REIT. The Simon brand was displayed prominently in all its malls and promoted through both hard-copy media and online. The idea was that if Simon malls were consistent enough in quality and service, then perhaps consumers would attach some value to the brand and choose Simon malls over other, more convenient malls when traveling. Moreover, brand recognition may work synergistically with geographical scope to build franchise value. For example, a tenant in a Post Property apartment decides to move to another city and needs to break her lease. Rather than charging the tenant a deposit penalty, the property manager allows the tenant to break the lease provided she moves into another Post Property in her new city. High-quality service and consistent amenities may help such a REIT keep tenants longer and experience lower overall vacancy rates. So while the strategy of building franchise value remains somewhat unproven in the REIT industry, there is some logic to support the idea of trying to build franchise value when combined with economies of scale and geographic scope.

Power in the Space Market The last major strategic consideration for REITs is the idea of a REIT using its size and access to capital to corner a local market for a specific property type by concentrating its space ownership within a few geographically confined space markets. Rental markets are often sufficiently segmented to enable such concentrated ownership to impart some market power to the dominant space owner, at least in the short run. A dominant local presence may also keep new development out. However, there is no convincing statistical evidence for the effectiveness of such strategies, and it is rare to find geographically constrained space markets that are so unique that they cannot face substantial competition from substitute locations, at least over the medium-term horizon in which new construction can occur. Nevertheless, large REITs do have some ability to pursue a market power strategy, and such a strategy probably does offer some potential profit. With the increasing consolidation in the industry, rental market power may be becoming more of a reality in some locations and property sectors.

advantages in terms of revenues and costs for a sample of European REITs. They do not find scale advantages in debt and overall capital costs. More recently, and despite the overall growth in average REIT size, Highfield et al. (2021) still find that larger REITs do better than smaller ones, both in revenue and in costs.

²⁹ Although liquid REITs do attract more institutional ownership (Cheung et al., 2015), Ambrose et al. (2019) do not find evidence of scale advantages in cost of debt and overall capital costs.

26.4 REIT INVESTOR CONSIDERATIONS

Even though REITs are listed stocks, many institutional investors do not manage their REIT exposure as part of their equity portfolio but as part of their overall real estate exposure. A key question for these investors is how to allocate their capital between the public and the private real estate markets.

Choosing between public and private real estate investment. An important question all real estate investors face is how to place their real estate investment capital: in the public equity market via REITs or in the private property market (directly or indirectly) via private market vehicles. While there are a variety of considerations in addressing this choice, the major trade-offs were touched on in Chapter 1, and you can understand them in more depth based on the portfolio theory presented in Chapter 22. REITs provide more liquidity than private market investments for the average investor, but they also provide less diversification within a wealth portfolio that is dominated by stocks and bonds, at least over short to medium investment horizons. This is because, on a contemporaneous basis, REIT returns tend to be more positively correlated with stock and bond returns than are private property market returns. However, a paper by Hoesli and Oikarinen (2021) looks at the behavior of REIT and private real estate returns for six countries and finds that these returns have similar characteristics and react similarly to economic shocks. This suggests that portfolio consideration of return and diversification are not very important when deciding how to allocate the real estate portfolio across REITs and private real estate.

The efficiency of the stock market can tend to give REIT investors more protection against making foolish mistakes, and therefore may require less intensive or less sophisticated and specialized due diligence and investment management, assuming investors use common sense and basic investment prudence. Thus, REITs are a logical route for small, nonspecialized investors who want some real estate exposure in their portfolio. However, Andonov et al. (2013, 2015) investigate the real estate investment approach of about 1000 pension funds worldwide and show that smaller funds investing in real estate are more likely to do so in private than in public markets. They also show that larger, more sophisticated investors, who will typically have the ability to invest in both REIT and private market investment vehicles simultaneously, are indeed more likely to combine these approaches.

Besides know-how, investment cost is an important criterion in deciding whether to allocate capital to REITs or to private real estate. REIT portfolio management costs are typically much lower than those for private real estate vehicles. Carlo et al. (2021) investigate these costs globally and show that pension funds investing in private vehicles typically pay annual asset management fees of around 100 basis points (1.0 percent) of asset value, for relatively simple “core” portfolios of stabilized properties. U.S. funds pay 125 basis points, which is considerably more than the 75 basis points paid by European funds and the 55 basis points by Canadian funds. This contrasts greatly with REIT portfolios, for which these costs vary around 25 basis points, depending on the country, the degree to which the portfolio is managed actively, and the question of whether pension funds manage these portfolios in-house or through an external asset manager.

A final consideration in choosing between public and private real estate investment is the availability of the different property types in these two markets. The property type allocation of an investor who would track the NAREIT index—either passively or with some active tracking error—would be driven by the industry’s cumulative choices for certain property types. The resulting allocation is likely to differ from the market composition and/or from the investor’s preferences. Of course, the investor may re-balance her property type exposure by overweighting and underweighting certain types of REITs, but when very large institutional investors do this, it may create unfavorable market impact, bidding up prices for REITs in property types that get overweighted in their portfolios. Taking positions in private real estate may be a better way to counterbalance unbalanced property type exposure.

Related to this is the fact that REITs tend to be a bit more innovative in finding “new” property types than the private markets. Many of the current ‘alternative’ property types started out as fringe

products and were then gradually brought to institutional respectability in the REIT market. So, for investors looking for such innovative property types, the REIT market has more to offer than the private real estate markets.

REIT versus general stock returns. For most REIT investors, their REIT holdings are part of a mixed-asset portfolio in which other sectors of the stock market play an important role. Given that mixed-asset diversification is a key consideration in including real estate in the portfolio, it is of particular interest how REITs perform relative to the stock market, including how this performance may differ between “up” and “down” stock markets. While past behavior in this regard is not necessarily a foolproof guide to future behavior, some pretty strong empirical evidence has been compiled on this topic. Much of this evidence seems to present a fairly consistent picture between the “old REITs” of the pre-1992 world and the “new REITs” of the 1990s and later.

In general, REITs are low-beta stocks, but they are certainly not zero-beta. REITs have consistently shown higher correlations (and betas) with respect to small stocks than with respect to large stocks. Note that REIT volatility has tended to be near that of large-cap stocks, which implies a beta with respect to large-cap stocks approximately equal to the correlation coefficient. Sing et al. (2016) show that the changing patterns in betas and correlations are mainly attributable to the industrywide dynamics in REIT leverage.³⁰

26.5 CHAPTER SUMMARY

This chapter provides an introduction to a group of firms that link the stock market and the private property market, allowing investors to get real estate exposure through the stock market, and owners and managers of property portfolios to access the public equity markets. REITs are unique, whether your perspective is that of a stock market investor or a private property market investor. REITs, essentially listed real estate vehicles that pay zero corporate tax, were a 1960 U.S. invention with a slow start, with very limited initial adoption both within and outside of the U.S. The structure only got momentum in the 1990s in the U.S. and thereafter gained remarkable traction internationally. In most developed equity markets, but also in many emerging markets, REIT-like structures have now been adopted for the public trading of real estate assets. REITs are an effective and relatively cheap way for investors to gain real estate exposure, and to do so in a range of property types, varying from the more traditional (offices, retail, multifamily residential) to the niche (student housing, self-storage, healthcare).

Invariably, the privilege of a zero corporate tax comes with obligations and prohibitions. These differ across countries, but two obligations are nearly always present. The first is to pay out most or all of the annual earnings as dividends, and the second is that the asset portfolio and revenues are mostly based on real estate. Besides that, REITs are often limited in their capital structure and in the range of their operational activities such as property management and development. However, the extent to which that is the case differs strongly across countries, with some countries allowing these activities in taxable affiliate organizations.

This chapter looks at REIT valuation and does that in two ways: from a discounted cash flow perspective, and from the perspective of the combined value of the underlying property portfolio. Of course, the values of the underlying properties are themselves the result of the discounted cash flows of these properties. Outcomes of one method can be calibrated by applying the other. Although the values of the public and private real estate markets cannot diverge structurally and over the long run, we do observe extended periods in which prices in the public market are higher than those in the private market, and vice versa. The former periods are associated with REIT IPO activity, secondary

³⁰ Case et al. (2010) provide evidence on REIT-stock correlations and betas for the 1972–2008 period, while Sing et al. (2016) extend that to 2013.

equity offerings, and REIT growth generally, while de-listings are more common in the latter. These pricing differences are partly due to the fact that REIT values adjust more speedily to changes in overall market conditions than private real estate markets, both in the down- and in the upswing.

KEY TERMS

- Asset management pure plays
- Look-through provision
- Five-or-fewer rule
- Taxable REIT subsidiary (TRS)
- Initial public offering (IPO)
- GAAP Net Income
- Funds from operations (FFO)
- Earnings before interest, taxes, depreciation, and amortization (EBITDA)
- Equity-before-tax cash flow (EBTCF)
- Adjusted funds from operations (AFFO)
- Funds, or cash, available for distribution (FAD or CAD)
- Gordon growth model
- Plowback ratio
- Dividend payout ratio
- Assets in place
- Same-store growth
- Growth opportunities
- Average equity cost of capital
- Net asset value (NAV)
- Growth stocks
- Income stocks
- Value stocks
- Public/private arbitrage
- Vertical integration

STUDY QUESTIONS

Conceptual Questions

- 26.1. What is the original purpose for which Congress enacted the REIT enabling legislation in 1960?
- 26.2. Why are REITs exempt from corporate income taxes?
- 26.3. What are the restrictions on REIT annual gross revenue sources and taxable income distributions necessary for a REIT to maintain its exemption from corporate income tax?
- 26.4. What are the two fundamental ways to think about the value of a REIT?
- 26.5. Define both “funds from operations (FFO)” and “adjusted FFO (AFFO),” carefully detailing differences between the two. Explain why these income measures are often used instead of GAAP Net Income to quantify the income-producing ability of a REIT.
- 26.6. Explain the Gordon growth model of stock market equity valuation. What are the major mathematical simplifying assumptions in this model?
- 26.7. What is the NAV of a REIT, and how is it estimated?
- 26.8. What is meant by REIT “public/private arbitrage”? How can REITs in principle create positive NPV for pre-existing shareholders when the share price to NAV premium is

positive ($P/NAV > 1$, private market values properties less than the stock market does)? How can REITs in principle create positive NPV for pre-existing shareholders when the share price to NAV premium is negative ($P/NAV < 1$, private market values properties more than the stock market does)?

26.9. Describe three of the six major REIT management strategy considerations or objectives (other than financing strategy).

26.10. Outline the advantages and disadvantages of investment in public REITs compared to direct real estate investment from investors' perspective, as you see them.

Quantitative Problems

26.11. Middlepoint Industrial Property Trust reports the following financial information, on a per-share basis:

GAAP Net Income [Earnings per Share (EPS)] = \$4.00

Depreciation and Amortization (GAAP) = \$3.00

Dividend Paid = \$4.75

Debt = \$50, with an average interest rate = 7%

Estimated share price premium to net asset value = 10%, based on an estimated private property market cap rate of 8% (weighted average firm level cap rate).

- Determine the minimum dividend the firm must pay to shareholders to be a REIT.
- Explain how the REIT can pay a dividend in excess of earnings per share. Calculate funds from operations (FFO) as part of your answer.
- Determine the aggregate NOI on Middlepoint's properties.
- Determine the current public market price of a share of Middlepoint stock.

26.12. Colonial Apartment REIT reports the following financial information on a per-share basis:

GAAP Net Income [Earnings per Share (EPS)] = \$4.00

GAAP Depreciation and Amortization = \$3.00

Reserves for property capital expenditures = \$0.80

Dividend paid = \$4.50

Debt = \$50, with an average interest rate on debt = 7%

AFFO yield of 6.9% based on current share price

Estimated weighted average firm-level private property market cap rate of 8%

- Determine Colonial's EBITDA.
- Determine the current (public) price of a share of Colonial REIT common stock.
- Determine the current share price premium (or discount) to NAV for the REIT.

26.13. Blackstone Properties, a national office REIT, reported GAAP Net Income of \$78,806,000 in 2005. It also reported the following income statement items:

Real estate depreciation: \$147,746,000

Preferred stock dividends: \$24,468,000

Allocations to OP unit holders: \$26,983,000

Net gains from sale of real estate assets: \$2,058,000 Net loss from extraordinary items: \$13,786,000

- What were Blackstone's funds from operations (FFO) for 2005?
- Given that Blackstone reported a 2005 net straight-line rental adjustment loss of \$14,619,000 and capital expenditures of \$27,500,000, what were Blackstone's adjusted funds from operations (AFFO)?
- What was Blackstone's 2005 AFFO/GAAP Net Income ratio?
- Given that the REIT paid out \$143,826,000 in common stock dividends in 2005, what was its dividend/FFO ratio?

- e. What was Blackstone's plowback ratio based on its AFFO (and deducting distributions to preferred shareholders and OP unit holders as well as common dividends)?
- f. By what percentage did Blackstone exceed the 90% minimum earnings payout requirement (based on GAAP Net Income)?

26.14. On January 31, 2005, National Growth Properties, a major shopping mall REIT, was trading at a share price of \$28.50, at which price its then-current (year-2005) annual dividend of \$2.04/share provided a yield of 7.2%. At that time, several analysts who followed National Growth closely were predicting future growth in earnings at a rate of nearly 10% per year for the subsequent five years, although same-store rental growth during the preceding year had been only 5.3%.

- a. If the analysts' earnings growth rate expectations accurately reflected the stock market's long-run average growth expectations for National Growth's dividends as of January 31, 2005, then, based on the GGM, what was the stock market's implied long-run average required expected rate of total return for investment in National Growth equity?
- b. If the market's long-run growth expectations were better reflected by the previous year's same-store rental growth rate, what was the market's required total expected return?
- c. With the 5.3% growth expectation assumption of question (b), what is the market's implied required *ex-ante* risk premium for National Growth equity, given that T-bills were yielding about 5.5% at that time?
- d. If the required total return expectation you calculated in question (b) was indeed the market's required return for National Growth, but the most accurate longterm dividend growth rate expectation for National Growth was actually halfway between the 5.3% same-store rental growth rate and the analysts' 10% earnings growth rate prediction, then what was the extent of the market's "underpricing" of National Growth stock on January 31, 2005?

26.15. The Wonderbloom Capital, a publicly traded REIT, has an expected total return to equity of 13%, average interest rate on its debt of 7.5%, and a debt/total asset value ratio of 40%.

- a. What is Wonderbloom's equity average cost of capital?
- b. What is Wonderbloom's firm-level overall average cost of capital?

26.16. Bob & Sue Realty (BSR) is a publicly traded REIT that has no debt and a current dividend yield of 8%, with a current share price/earnings multiple of 12.5. The current consensus expectation among stock analysts who follow BSR is that BSR can provide a long-term average growth rate in its dividends per share of 5% per year.

- a. What is BSR's plowback ratio (i.e., what proportion of its earnings does it retain and not pay as dividends)?
- b. Assuming the stock market agrees with these analysts' expectations, what is BSR's firm-level average cost of capital? (Hint: As BSR has no debt, you can use the GGM directly to answer this question.)

26.17. The Rentleg Distribution Center is a warehouse complex near the Cincinnati Airport in Northern Kentucky, in a market where such buildings currently sell at 10% cap rates (net cash flow/property value), with 1.0% expected long-run average annual growth (in both value and cash flow). This property has an initial net cash flow of \$2,500,000 per year. Both the Grump REIT and BSR (from Questions 23.21 and 23.22) are considering bidding to buy the Rentleg Center.

- a. Ignoring possible differential valuation between the stock and property markets, what is the (marginal) opportunity cost of capital for acquisition of the Rentleg Center by either REIT (i.e., what is the *discount rate* relevant for a DCF valuation of the Rentleg Center on the part of either one of the two REITs)?

- b. What is the maximum price Grump can offer for the Rentleg Center without its share price being diluted?
- c. What is the maximum price BSR can offer for the Rentleg Center without its share price being diluted?
- d. If BSR is able to purchase the Rentleg Center for \$24,000,000, what will be the change in the (aggregate) value of BSR's equity as a result of this transaction (assuming the stock market had not already factored such an expected purchase into its valuation of BSR's shares)?

REFERENCES AND ADDITIONAL READING

Ambrose, B., S. Erlich, W. Hughes and S. Wachter. 2000. REIT Economies of Scale: Fact or Fiction?, *Journal of Real Estate Finance and Economics* 20: 211–224.

Ambrose, B., M. Highfield and P. Linneman. 2005. Real Estate and Economies of Scale: The Case of REITs, *Real Estate Economics* 33: 323–350.

Ambrose, B., F. Fuerst, N. Mansley and Z. Wang. 2019. Size Effects and Economies of Scale in European Real Estate Companies, *Global Finance Journal* 42(C).

Andonov, A., P. Eichholtz and N. Kok. 2013. A Global Perspective on Pension Investments in Real Estate, *Journal of Portfolio Management* 29(5): 32–42.

Andonov, A., P. Eichholtz and N. Kok. 2015. Intermediated Investment Management in Private Markets: Evidence from Pension Fund Investments in Real Estate, *Journal of Financial Markets* 22: 73–103.

Boudry, W., J. Kallberg and C. Liu. 2010. An Analysis of REIT Security Issuance Decisions, *Real Estate Economics* 38: 91–120.

Boudry, W. 2011. An Examination of REIT Dividend Payout Policy, *Real Estate Economics* 39(4): 601–634.

Bradley, M., D. Capozza and P. Seguin. 1998. Dividend Policy and Cash Flow Uncertainty, *Real Estate Economics* 26: 555–580.

Capozza, D. and P. Seguin. 1998. Managerial Style and Value: The REIT Evidence, *Real Estate Economics* 26(1): 131–150.

Capozza, D. and P. Seguin. 1999. Focus, Transparency and Value: The REIT Evidence, *Real Estate Economics* 27(4): 587–620.

Carlo, A., P. Eichholtz and N. Kok. 2021. Three Decades of Global Institutional Investment in Commercial Real Estate, *Journal of Portfolio Management* 47(10): 25–40.

Case, B., W. Hardin and Z. Wu. 2012. REIT Dividend Policies and Dividend Announcement Effects during the 2008–2009 Liquidity Crisis, *Real Estate Economics* 40(3): 387–421.

Cheung, W., R. Chung and S. Fung. 2015. The Effects of Stock Liquidity on Firm Value and Corporate Governance: Endogeneity and the REIT Experiment, *Journal of Corporate Finance* 35: 211–231.

Demirci, I., P. Eichholtz and E. Yönder. 2018. Corporate Diversification and the Cost of Capital: The Case of REIT Mortgages and Bank Loans, *Journal of Real Estate Finance and Economics* 61(3): 316–368.

Derwall, J., J. Huij, D. Brounen and W. Marquering. 2009. REIT Momentum and the Performance of Real Estate Mutual Funds, *Financial Analysts Journal* 65(5): 24–34.

Dogan, Y., C. Ghosh and M. Petrova. 2019. On the Determinants of REIT Capital Structure: Evidence from Around the World, *Journal of Real Estate Finance and Economics* 59: 295–328.

Ertugrul, M. and E. Giambona. 2011. Property Segment and REIT Capital Structure, *Journal of Real Estate Finance and Economics* 43: 505–526.

European Public Real Estate Association (EPRA). 2022. *EPRA Global REIT Survey 2022: A comparison of the major REIT regimes around the world*.

Fama, E. and K. French. 1992. The Cross-Section of Expected Stock Returns, *Journal of Finance* 47: 427–465.

Feng, Z. and P. Liu. 2023. Introducing “Focused Firms”: Implications from REIT Prime Operating Revenue, *Journal of Real Estate Finance and Economics* 67, 545–578.

Gordon, J. and E. Shapiro. 1956. Capital Equipment Analysis: The Required Rate of Profit, *Management Science* 3: 102–110.

Gyourko, J. and E. Nelling. 1996. Systematic Risk and Diversification in the Equity REIT Market, *Real Estate Economics* 24(4): 493–515.

Highfield, M., L. Shen and T. Springer. 2021. Economies of Scale and the Operating Efficiency of REITs: A Revisit, *Journal of Real Estate Finance and Economics* 62: 108–138.

Hoesli, M. and E. Oikarinen. 2021. Does Listed Real Estate Behave like Direct Real Estate? Updated and Broader Evidence, *Applied Economics* 53(26): 3023–3042.

Kim, D. and J.A. Wiley. 2018. NAV Premiums and REIT Property Transactions, *Real Estate Economics* 47(1): 138–177.

Peterson, J. and C. Hsieh. 1997. Do Common Risk Factors in the Returns on Stocks and Bonds Explain Returns on REITs?, *Real Estate Economics*: 25: 321-345.

Ro, S. and A. Ziobrowski. 2011. Does Focus Really Matter: Specialized versus Diversified REITs, *Journal of Real Estate Finance and Economics*: 50–72.

Sing, T.F., I. Tsai and M. Chen. 2016. Time-Varying Betas of US REITs from 1972 to 2012, *Journal of Real Estate Finance and Economics* 42: 68–83.

27 International Real Estate Investments

Markets, Strategies, and Implementation

CHAPTER OUTLINE

- 27.1 There's a Big World Out There
 - 27.1.1 The Global Real Estate Capital Market
 - 27.1.2 The Institutions of Globalization
- 27.2 Rationales for Going International
 - 27.2.1 Return Opportunities
 - 27.2.2 International Diversification
 - 27.2.3 Managerial Considerations
- 27.3 Obstacles to International Property Investment
- 27.4 Developing and Implementing International Real Estate Strategies
 - 27.4.1 Public Real Estate Investment
 - 27.4.2 Determining Country Allocation
 - 27.4.3 The "Home Market" Concept
- 27.5 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The rationale of international real estate investment.
- The major obstacles and disadvantages to investing in real estate in foreign markets.
- The nature of international institutional investment portfolios and capital flows and the practical considerations shaping them.
- The nature of the institutions enabling a successful international real estate strategy.
- Risk management strategies associated with successful international real estate investment.

Real estate investment was traditionally a very local business with investors focused on nearby markets they knew well. Regional diversification has spread to take advantage of different growth rates while, at the same time, capital markets have become more international. Today, cross-border capital flows have increased in all asset markets. Real estate is following suit and is now beyond the tipping point, where more international investment leads to new investment products and to the emergence of supporting institutions that, in turn, facilitate yet more international capital flows. The development feeds on itself.

The preceding chapters have looked at real estate markets on a regional or national level, with the last chapter looking at the development of the global REIT market. REITs are well placed as conduits for international real estate investment, which makes this a logical place in the book to look at international property investment in general. This chapter will begin by looking at the global property investment universe and at international capital flows, and will discuss the institutions aiding the

internationalization of the global property markets. After that, we will consider the key arguments for international real estate investment: return opportunities abroad, diversification, and the export of portfolio management and development skills. We will also examine the obstacles and problems associated with international property investment, such as additional costs and risks. The key issue here is overcoming the informational disadvantage that distant property investors have relative to their local peers.

Finally, we will synthesize these issues and review viable portfolio strategies for international property investors, based on a combination of direct property investment and the use of the global REIT market. One of the key insights is that different types of investors have different optimal international investment strategies. Especially important in that regard is the question whether one is an intermediate investor, whose shareholders can diversify risk themselves, or an end-investor, who does not have diversifying shareholders. The chapter ends with suggestions for implementing these strategies, discussing ways to determine country allocation, global indexing strategies, and currency risk management.

27.1 THERE'S A BIG WORLD OUT THERE

The global real estate market is growing, as can be measured by looking at invested capital and capital flows, as well as the size of the market and market segments. This parallels the growth of international corporations with real estate needs around the globe. Besides that, the institutions of globalization are emerging in real estate. As in any sector, there is a chicken-and-egg situation where international capital flows are aided by transparency-enhancing market institutions, but for these institutions to emerge and be sustained, the capital flows must emerge. Beyond a certain threshold, the two feed on each other, and the emergence of institutions of international transparency is both a sign of and a catalyst for advancing globalization.

27.1.1 THE GLOBAL REAL ESTATE CAPITAL MARKET

In general, investors favor their own local markets above asset markets far away. This so-called home bias has been well documented for stock and bond markets, and real estate markets are no different in that respect. The mantra “location, location, location” suggests that home bias could even be more salient for real estate than for other asset markets since local factors drive both supply and demand. Historically, real estate investors were reluctant to expand internationally, and that held particularly true for U.S. investors: Of all foreign direct investments made from the U.S., average real estate investments have been a mere 0.3 percent in the decades through 1995.¹

But internationalization is underway, and the global market has become bigger and more transparent. The three main destinations for these international capital flows are China, the United Kingdom, and the United States, in that order, but where this capital largely flows into standing investments in the latter two countries, the bulk of the capital flowing to China is destined for new property development. Other major recipients are Germany, Japan, and France.

Interestingly, the degree to which these international capital flows are truly global or rather more regional varies by region. The United Kingdom and the United States attract a truly global flow, while the capital going into China is almost exclusively regional in nature. For capital destinations in continental Europe, and for Japan and Australia, global and regional sourcing balance out.

By far the largest exporter of real estate capital is the United States, investing predominantly outside its region. Singapore, Canada, Hong Kong, and the United Kingdom are also large exporters of capital into foreign real estate markets, but they have a strong regional focus.

It is also clear that international capital markets are integrating, especially within Europe. Cross-border real estate capital flows among European countries are expanding, and they tend to grow faster than domestic investments. However, many investors stay local, illustrated by the behavior

¹ For several countries between the 1960s and 1995, Lapier (1998) shows that international investments do not go up as a share of total private real estate investments, and stay at very low levels, suggesting that the home bias was much higher for real estate than for other assets.

of listed property companies. In the global universe of property companies, as measured by Global Property Research (GPR), the percentage of internationally diversified companies has been falling over time, slowly decreasing from about 10 percent in 2007 to 6 percent in 2024. This illustrates that real estate companies have largely remained local. International capital market integration seems to go along regional not global pathways. This notion is in line with existing empirical research, which also suggests that real estate markets are integrating on a regional basis.²

27.1.2 THE INSTITUTIONS OF GLOBALIZATION

International capital market integration requires institutional help, both through the emergence of investment products catering to international investors, and through reliable information sources that make markets more transparent. The development of the public real estate markets, both on the equity and on the debt side, makes it much easier for property investors to allocate significant amounts of capital outside of their home market. The previous chapter showed that the global property share market has been growing in the last three decades, helped by the proliferation of tax pass-through structures all over the world.

One can also view the global market from a debt point of view. Globally, the market for Commercial Mortgage-Backed Securities (CMBS), which seemed like useful vehicles for real estate debt investors to step into foreign debt markets, had been growing until the 2008/09 crisis, but has been more or less closed down since then. That crisis has also demonstrated the disadvantages of property market globalization: securitized debt markets have been an important transmission channel of the crisis from one country's banks and investors to the next.

Emerging globalization seems to go hand in hand with a trend for global property markets to become more public. This, in turn, facilitates more internationalization. It is far easier for an investor in Amsterdam or London, for example, to invest in U.S. property by buying shares of listed property companies on the exchange than to travel there physically, open an office and set up business as a direct investor. So, while the emergence and growth of the securitized market has not by itself made the market international, it has been a conduit and enabling factor for internationalization.

Besides the increasing publicness of real estate markets, there are also other new institutions, important for investors wishing to make direct investments in private real estate. In a private market, the relevant knowledge the real estate investor needs to be successful and perform well remains largely local in nature. In the past, this knowledge was not shared as there were no associations, agencies or advisors willing to share it, but these have now been created.

This results in the establishment of standards in information quality, governance, and professionalism, both via professional organizations and for-profit firms. National and international research consultants, data providers, and brokers increasingly supply local data and can help investors build up local knowledge. These data include performance benchmarks, up-to-date market information on property transactions, rents, vacancies, and yields, and information regarding rent contracts up to the individual property level. The quality as well as availability of these data are constantly improving. However, this holds true mainly for the mature markets, such as North America, Europe, Japan and Australia. This transparency is less strong for important emerging property markets like those of China, India, Turkey, and Brazil.

International transparency is demonstrated by the evolution of Jones Lang LaSalle's transparency index, a global real estate indicator that measures and ranks countries in order of transparency on a range of performance-relevant aspects. Since its inception in 1999, this indicator has shown ever-increasing transparency, and until the Global Financial Crisis of 2008–2009, no country had decreased its transparency. In the last decade and a half, however, transparency trajectories are diverging. The most transparent countries continue to improve, particularly on issues such as climate action, technology adoption, and anti-money laundering measures. On the other hand, in the middle

² See Eichholtz et al. (1998), who provide evidence that property markets are driven by continental not global market factors. Gerlach, Wilson and Zurbruegg (2006) show that the Asian crisis of the late 1990s resulted in stronger integration among Asian property markets. More evidence for regional integration is provided by Gallo and Zhang (2010) and Zhou (2011).

and bottom of the list, countries are decreasing in transparency, after having made important steps forward in the past. About a quarter of countries have seen their markets become less transparent during the last decade. Between 2016 and 2022, 15 countries have even dropped out of Jones Lang Lasalle's universe. Transparency clearly remains a problem in these markets.³

To sum up these developments, the long-term trend is for real estate markets to become more international, aided by institutions like growing public markets, performance indices and reliable and accessible market information. This leads to the question of why these developments are taking place, and what investors should do about them. Should they go international, and if so, how? In the next section, we consider the rationale for and obstacles to going international.

TEXT BOX 27.1 INSTITUTIONS OF INTERNATIONAL TRANSPARENCY

The internationalization of real estate capital flows has led to the emergence of supporting institutions. These institutions, both industry associations and for-profit companies, create common international standards and definitions, compare best practices and provide market information. These efforts result in increasing transparency, which facilitates yet more growth in international real estate capital flows. The list below provides information regarding some of the most useful of these institutions. The list is by no means exhaustive.

- www.aprea.asia: the Asia Pacific Real Assets Association, the leading industry association for the real estate sector in Asia.
- www.epra.com: the European Public Real Estate Association. The website provides indices and research reports, plus links to members' websites.
- www.nareit.com: the National Association of Real Estate Investments Trusts.
- www.ncreif.org: the National Council of Real Estate Investment Fiduciaries. The website provides lots of data regarding the private U.S. real estate markets.
- www.inrev.org: the website of INREV, the European Association of Investors in non-Listed Real Estate Vehicles. Provides useful information for members.
- www.naidirect.com: up to date direct property market reports globally.
- www.msci.com: site of the world market leader in private property performance indices, the data that was formerly provided by IPD and RCA; provides performance information for real estate markets in many countries.
- www.us.jll.com / www.jll.co.uk / etc: international real estate advisors. Provides research reports and updates of the global real estate transparency index.
- www.transparency.org: Transparency International, the main international organization measuring corruption; produces an annual corruption report.
- www.cbre.com: The website of global real estate advisor CBRE, provides extensive up-to-date market studies on its website, which are freely accessible.
- www.colliers.com: The website of Colliers International, provides market reports for registered visitors.
- www.costar.com: The website of CoStar, a leading building-level data provider.
- www.globalpropertyresearch.com: The website of Global Property Research, a maker of international property share performance indices.

27.2 RATIONALES FOR GOING INTERNATIONAL

There are strong arguments in favor of international real estate investment. First, there may be good investment opportunities offering better returns outside the home country, especially if the home

³ Lieser and Groh (2011) analyze the determinants of international commercial real estate investment in 47 countries and show that (lack of) real estate market transparency plays a significant role, besides issues like economic growth, urbanization, and demographics.

country has a well-developed, mature property market where the rate of urbanization is already high, and little growth opportunities remain. Another rationale is that international investment provides diversification benefits because markets do not move in synchronized ways. The third reason could be to export superior portfolio management or development expertise, especially in emerging markets, where such expertise may be in demand.

27.2.1 RETURN OPPORTUNITIES

Prospects of better return opportunities than in the home country come in two varieties. First, return opportunities can be structural, relating to economic development, relative capital scarcity, and demographics. Secondly, return opportunities abroad may be of a cyclical or tactical nature. The investor's home market may be at the top of the property cycle, with a boost to returns beckoning from switching to another market at the bottom of the cycle.

Structural return opportunities

In economically mature markets, economic growth will be relatively low, and so will growth in property demand. But in emerging markets, rapid economic growth will spur growth in commercial property demand. Similarly, markets may be mature or immature in terms of urbanization and demographics. These factors are not always positively correlated. For example, Eastern Europe is economically an emerging region but demographically mature and with a high urbanization rate. China is an economically booming market, but already reached its population peak in 2022. In a country such as India, on the other hand, the peak in population is still decades away, the urbanization rate is still quite low, and economic growth prospects are good.

Regarding the first type of structural return opportunity abroad, economic development is obviously not evenly spread across the globe, with the differences partly of a structural nature. Economic growth in mature markets in Europe and Japan is relatively weak, and that implies slow growth in property demand. By contrast, emerging markets in Asia, Central Europe, and the Americas have (much) higher economic growth rates.

Even while China's economic growth has slowed down, it is still higher than in most developed economies. India has shown average growth of over 6 percent since 2000 and looks likely to continue on that path. Central Europe has shown average growth surpassing 3 percent since 2000. These regions will continue catching up with the mature economies. Until they do, they are likely to sustain high growth numbers, even though the growth path will probably not be very smooth.

Economic growth goes hand in hand with property demand. A characteristic of emerging economies is rapid urbanization. China, for example, is creating cities in the millions at a startling pace, housing the rural migrants partaking in the economic boom. Meanwhile, those already living in the cities are becoming richer and consume more space in the form of better housing. India and most African countries still have a very long way to go in that regard, creating huge investment opportunities for decades to come.

Developing societies also need industrial and office properties to accommodate workers. Cities attract migration from the countryside because there are jobs available, initially in manufacturing and distribution, and this requires industrial space such as warehouses and factories, as well as housing for the new urban population. Subsequently, service industries start growing rapidly, requiring office space. And as societies become more prosperous, populations adapt similar spending patterns to the West, creating demand for shopping and leisure projects.

This growing property demand leads to large structural capital needs, as housing people at work or home in high-growth urban areas absorbs huge amounts of capital. By contrast, many of the mature economies have mature capital markets, with large institutional investors and lots of capital chasing a limited supply of investable assets. This is nothing new: in the eighteenth century, the Dutch had a small economy with a big capital market, so their capital went abroad and financed the emerging American economy. Capital markets shift capital from places where it is in surplus to places where it

is in need. This probably explains why Dutch and British property investors traditionally have been relatively international: they have had large capital markets relative to their own real estate markets.

Demographic developments

The other major structural factor that could influence returns, vary widely across countries. Populations in some countries keep on growing, while in others, they are levelling off or already shrinking. Real estate markets provide the space that people need—working space, living space and recreational space, and these needs result largely from population size and composition, that is, from demographics. To get an idea of future property demand, we need to know how many people there will be as a whole and in different age categories.

Moreover, different property types give investors exposure to different aspects of demographic change. Demand for office space depends on the number of people of working age, as these will populate the offices. Demand for retail space depends partly on the total population (because everyone needs clothes and food), and on the population composition, as some age groups consume more than others do. Purchasing power is a key issue: there are far fewer shops per 1,000 inhabitants in Eastern than in Western Europe, but if purchasing power rises to Western European levels, shopping space may also reach these levels. Housing demand largely depends on household formation, which is a function of population and average household size: even if the population shrinks, a decreasing average household size can still increase the number of households, and therefore the demand for housing units.

However, the exact relationship between demographics, property demand, and property prices remains unclear. For example, the end of the baby boom in the U.S. is a rather non-dramatic event compared to demographic developments in other countries, since it does not imply a shrinking population but merely a slowing down of population growth: the United Nations does not predict an end to U.S. population growth for decades.⁴ But in some important markets outside the U.S., the population is indeed shrinking, and for real estate markets, with their inelastic supply, there is a fundamental difference between slow growth and a population that actually shrinks.⁵

Employment is the key demand driver in office market models, and over time, employment fluctuates with the economic cycle. However, job creation matters for office demand only as long as there are people to take the jobs. Exhibit 27-1 shows index numbers of the labor population—the number of people aged 25–64—in Asia, North America, Europe, and the world, starting at 100 in 1970. Asia shows the fastest historic growth but will top out in about 2045 and shrink thereafter. North America has had strong growth and keeps on growing through 2040, after which the labor population will enter a steadier state. In Europe, the situation is totally different, with the labor population already having topped out in 2015 and projected to fall below the 1970 level by 2060, though there will be a structural difference as virtually all women will work in 2060 compared with hardly any in 1970. This also means there is some potential for growth in the active labor force left in countries where few women currently participate in the workforce.

Among the mature markets, in the immigration-based group of economies such as Australia and the U.S., labor population growth will continue, albeit at a slower pace. But in mature economies elsewhere, Germany and Japan are already well past their peaks; Italy and the Netherlands have reached their peaks and Singapore and Spain are close to it. So, besides the immigration-based countries, there are two categories of mature markets. The “not-so-bad” group, like France, Sweden, the United Kingdom, and Singapore, where labor population is projected to be stable or just slightly falling, and the shrinking labor force group, including countries like Germany, Italy, Japan, and South Korea, which will all see their labor populations go down with more than 25 percent until 2070.

One solution could be to shift the property allocation to emerging markets. Here, we can see that markets like Brazil, India and Turkey look much more promising regarding demographics. But in

⁴ Francke and Korevaar (2022) look at very long-term data and find that an increase in the birth rate predictably increases house prices 25–30 years later, due to housing demand effects.

⁵ The key theoretical paper regarding this asymmetry is Glaeser and Gyourko (2005), who show that decreasing population may lead to strong declines in property prices. Using very long-term rental data, Eichholtz, Korevaar and Lindenthal (2023) empirically show that the effect of population change on rents is far greater on the downside than on the upside.

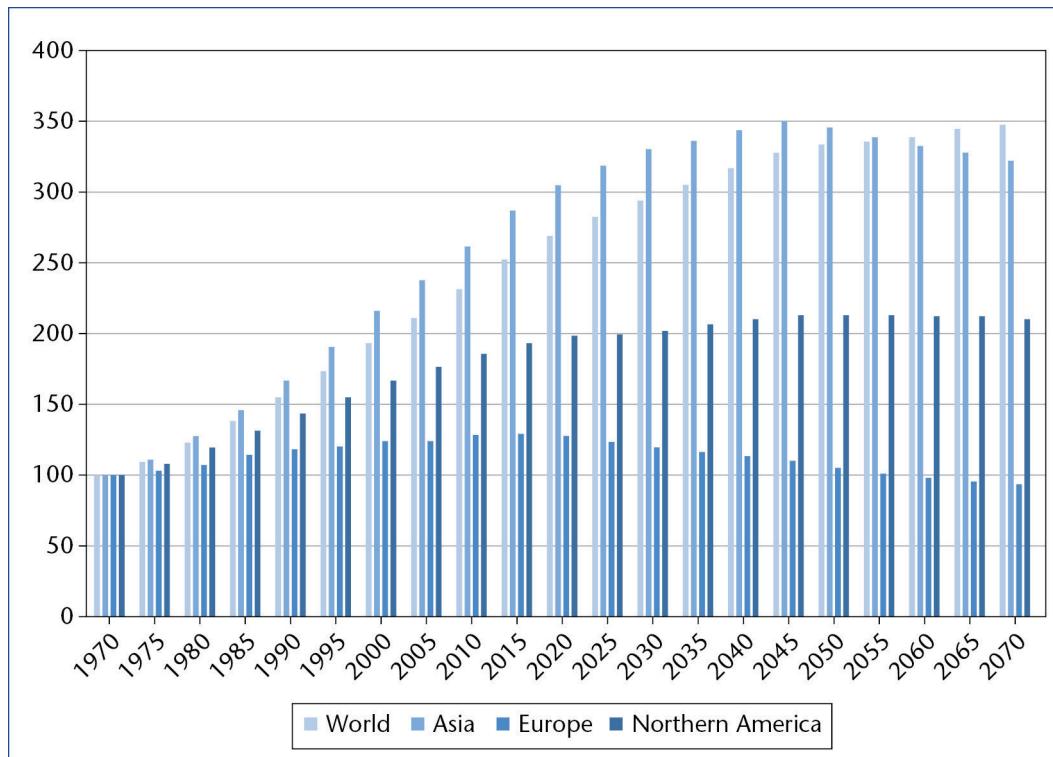


EXHIBIT 27-1 International Labor Force Statistics and Projections (age group 25–64).

Notes: The numbers have been indexed with 1970 as base year. In 1970 there were 1.44 billion persons aged between 25 and 64. Of this total Europe housed 318 million, North America housed 99 million and Asia housed 790 million persons.

Source: United Nations World Population Prospects Database.

Central Europe, growth is topping out or already falling: these countries may be economically emerging, but they are very mature demographically.

The next question is whether the working population even needs offices. Here, the answer has traditionally been yes because employment growth in mature economies was mostly in services. While this looks set to continue, future service employment may require (much) less office space than in the past. The organization of work in services is changing fundamentally due to the revolution in communication and information technology. This was already going on, but the global COVID pandemic has accelerated this trend. People work more from home and so-called “third places” including coffee shops, airport lounges and libraries, essentially anywhere they happen to be or want to be.⁶ However, the degree to which this happens varies a lot across countries and seems to affect the mature markets of Europe and the U.S. more than Asia, for example.

For mature office markets, demographics and communication technology make for structural gloom. But there may be some offsetting forces. First, structurally tight labor markets may induce more immigration, later retirement, and stricter social benefits policies. However, these developments are much less easy to predict than the demographic growth of the past. Another key

⁶ Gupta, Mittal and Van Nieuwerburgh (2024) provide a very thorough analysis of the effects of work from home on office demand.

factor will be migration within countries, with growing regions and cities pulling in labor from rural and stagnant areas. In real estate, competition between cities will grow with the increasing differences between fast-growth and stagnant cities. This will imply more uncertainty regarding the demand for space, and therefore more investment risk. Where the office market used to be driven by a constantly humming demographic engine, it must now rely on hard-to-predict political and social developments. Even if all possible mitigating factors do emerge, global office demand will grow less, due to less labor and the increasing use of alternative locations for work. While some places will continue to see growing demand for office space, that is not the case globally. New supply will have to slow down, or prices will fall, and developers as well as investors will feel this.

Macro-demand for retail space, on the other hand, is a function of population size and purchasing power. The number of people and their age composition determine the quantity of retail demand, and therefore of retail space. More people imply more retail purchasing, and different age groups have different consumption patterns. Meanwhile, purchasing power will determine quantity and quality. Net new space will only be needed if these variables grow, otherwise only replacement space will be required.

Exhibit 27-2 shows a strong variety in historic and expected population developments. Until 2070, the U.S. population is expected to grow approximately 15 percent, while a shrinking population is expected in some European countries—Germany, Italy, Spain, and those in Central Europe—and in China, Japan, and South Korea. In general, the numbers look better than those for the labor population, with expected population growth especially in emerging and immigration-based economies. Nevertheless, even in these countries, expected population growth for the next 50 years is much lower than it was for the past 50 years.

Moreover, purchasing power, the second quantitative driver of the demand for retail space is likely to continue growing. Simple extrapolation of real GDP growth since 1970 suggests that by 2040, even the relatively slow growers in Western Europe and North America will be substantially better off than in 2020, suggesting plenty of scope for growth in retail property. In the rest of the world, the potential for real GDP growth is (much) larger, which will translate into more demand for retail space. And as consumers become more affluent, they will demand a Western-style shopping experience, rather than the more informal retail arrangements that are still prevailing in many emerging economies.

But other than in the past, these developments will not translate automatically into consumer spending in physical shops. The global communications revolution breaks down the traditional relation between growth in population and affluence and the demand for retail space. In 2022, e-commerce accounted for nearly 20 percent of global retail purchases, and this percentage is growing steadily. At the same time, global retail sales in physical stores seem stuck at around 20 trillion U.S. dollars, suggesting that the global need for retail space will not grow. Of course, that does not hold for each country and region, but it does imply much bigger differences between regional markets for retail space. While cities in countries with strong urbanization will still need more physical retail space, the more mature countries will need (much) less of it in total, although new kinds of space requirements suggest we will still see development and redevelopment.

The final question is how demographic changes will affect residential properties. This market is driven by income trends, interest rates and planning rules in the short term, but ultimately by household formation, as every household needs at least one home. Household formation is a function of population growth and average household size. Households are generally still getting smaller, partly offsetting expected decreases in population in many countries. For example, while Germany's population has stopped growing, its number of households continues to increase, albeit at a slower rate than in the past.

For countries in which the population keeps increasing, the falling average family size is a factor behind even higher growth for housing demand. Especially in emerging economies, the potential for increased housing demand due to reduced family sizes is big, since economic development is very closely connected with lower fertility rates and smaller families. So even if population growth slows

EXHIBIT 27-2

Population (x 1000 persons) in International Perspective; 1970–2070

Country	Total Population	Cumulative Growth		Peak in
	2020	1970–2020	2020–2070	
A. Mature Economies				
Australia	25500	99.3%	45%	/
Canada	37742	76.6%	33%	/
France	65274	28.6%	2%	2046
Germany	83784	6.6%	-9%	2021
Italy	60462	13.0%	-23%	2017
Japan	126476	20.5%	-28%	2009
Netherlands	17135	31.8%	-4%	2035
Singapore	5850	182.3%	4%	2042
South Korea	51269	59.2%	-25%	2024
Spain	46755	53.8%	-19%	2020
United Kingdom	67886	22.2%	12%	/
United States	331003	77.3%	22%	/
B. Emerging Economies				
<i>Americas</i>				
Brazil	212559	123.5%	2%	2045
Chile	19116	95.4%	2%	2048
Mexico	128933	150.4%	21%	2062
<i>Asia-Pacific</i>				
China	1439324	73.9%	-13%	2031
India	84339	141.8%	15%	2059
Indonesia	273524	138.3%	23%	/
Malaysia	32366	199.6%	30%	/
<i>Europe</i>				
Czech Republic	10709	9.1%	-5%	2026
Hungary	9660	-6.8%	-21%	1980
Poland	37847	16.0%	-24%	1996
Russia	145934	12.1%	-11%	1993
Turkey	84339	141.8%	15%	2058
C. World	7794799	110.6%	34%	/

Source: United Nations World Population Prospects Database.

down due to that same lowering of the fertility rate, the growth in the number of households is set to continue for decades, ensuring ongoing demand growth for housing. For example, the average household size in Turkey has been going down in each census since 1975: it was four in 2008 and has fallen to 3.2 in 2022. If Turkey's household size were to go to the European average in the long run, this alone would lead to a doubling in the demand for housing units.

The other side of the demographic coin is that societies age, and this also seems to have an important impact on the demand for housing. Research suggests a strong and positive relation between human capital and the qualitative demand for housing, and it also suggests that this relation is stronger for older families.⁷ Since younger generations are better educated than their elders

⁷ See Eichholtz and Lindenthal (2014).

in almost every country, and especially so in emerging economies, our societies will increase their human capital as time progresses. This, together with the average age of the population, will create demand for higher-quality housing, which implies the need for a continuous adaptation of the housing stock.

Summing up demographic developments, we can say that property markets in many countries are turning into replacement markets, in which good quality assets will drive out bad quality, and competition among cities and regions will increase. At the same time, fundamental uncertainty in these property markets will grow as population growth comes to a halt. Because different property types are exposed to different aspects of the demographic tide, demographic changes will first affect office markets and only later, retail and housing markets. Europe and some countries in Asia are very mature demographically, while many emerging countries and countries with a tradition of immigration will follow these trends later. Thus, investors from demographically mature markets should make strategic allocations abroad, especially to demographically immature markets.

What does this imply for a property investor? That depends on the home market. For European investors, these numbers provide a clear rationale for investing internationally, and that also holds for investors from China, Japan and South Korea. For American and Australian investors, where the demographics at home look much better, it means they must be very careful where to invest to achieve a better return/risk profile and they should be aware that there is no strict one-to-one relationship between economically emerging markets and demographically emerging markets. The same holds for investors from other countries in which the demographic situation looks favorable. In short, for one group of investors, demographics provides a rationale for going international; for another group, it is something they must be wary of.

Cyclical return opportunities

The structural return opportunities abroad require a long-run perspective on foreign markets, but the real estate cycle may lead to more tactical opportunities. The existence of cycles is well documented, and at any given point in time, and for any property type, there will probably be markets around the world at every phase in the cycle. Looking at international cycle snapshots supports this notion.

In theory, an opportunistic office investor who feels that the home office market is at the top of the cycle can sell assets and go into foreign office markets close to the bottom of their cycle. Such a strategy of opportunistic international cycle surfing sounds attractive, but to put these cyclical movements into practical investment policies, they need to be predictable. It is not clear whether that is the case, and whether persistence in real estate return series allows profitable tactical allocation policies, especially after adjusting for transaction costs⁸. The available information on which these trading decisions must be made is always backward-looking, which may be suggestive of future directions, but does not really make the markets predictable. Moreover, the inherent illiquidity of real estate may prevent such tactically timed transactions.

The international dimension makes matters even more difficult. Opportunistic international property funds investing in, for example, Central and Eastern Europe and in Asian markets like India, have been popular with investors for a while, but generally have not seemed to be very successful in terms of performance. Predicting property markets is hard, and this illustrates that it is even more difficult to make good predictions for someone who is located far away. There is the option of paying a local expert to provide knowledge about the market, but it is questionable whether local players with genuine knowledge will pass this on, rather than keeping it for their own benefit. After all, foreign investors are easier to exploit than their local peers. Another option, especially when

⁸ Marcato and Key (2005) look at British private real estate return series and show that active momentum strategies based on time series information generated from these series are profitable, even on a net basis, but they assume instantaneous transactions. Brounen, Eichholtz, and Ling (2007) show that active property trading does not lead to significant outperformance. As far as we know, there is no evidence on international trading.

new to a market, is to find a local partner, where interests can be aligned. Such a strategy also works domestically when venturing into a new market or markets where regulatory hurdles are tricky and local contacts are helpful.

27.2.2 INTERNATIONAL DIVERSIFICATION

In considering cyclical rationales for going abroad, we have concluded that markets are not very predictable, and even if they are, acting on the predictions is very difficult, especially for foreign investors. But we do know that markets move in non-synchronous ways, and this can be used for diversification. Creating exposure to different segments of the property market provides diversification because business cycles are not synchronous but move out of phase. Countries also have different economic bases, with for example the Netherlands almost entirely a service economy, South Korea a very strong industrial economy, Canada more of a mining and agriculture economy, and the U.S. a big mixed economy. Exposure to different economies also provides diversification via exposure to different economic bases and thus reduces risk. Besides that, property supply is partly driven by local capital availability and the interest rate, as well as zoning rules, and these vary from country to country as well.

Diversification potential can be measured by looking at correlations. If the fundamentals driving property returns—like GDP growth, inflation, and interest rates—show low international correlations, this is likely to translate into weakly correlated property markets, and therefore strong diversification potential. To look at that, we have calculated correlations for these fundamentals across several important economies and property markets, based on annual data going back to the mid-1990s until 2022. The resulting international correlation matrices are depicted in Exhibit 27-3. These numbers look promising for international diversification potential, as the correlations are generally low. For example, the correlations between the Euro-area's GDP growth and that of the other economies are low, except for the U.K., the U.S. and Japan. The fundamentals suggest strong diversification potential, and that is clearly visible in the timing of property cycles across the continents. These show phase variability: there are markets to be found in almost any phase of the cycle whenever a snapshot is taken, and this holds for all property sectors.

Additionally, rental cash flows and their volatility are driven by local market institutions such as rental contracts. In the Netherlands, for example, rental contracts are linked to inflation, partially protecting an office landlord from inflation risk, whereas in other countries, that's not so. But there, retail rental contracts may be linked to store turnover, directly linking the investor to consumer confidence and consumer spending. As such practices differ across countries, this creates diversification. The international differences in rental contracts provide diversification benefits through different cash flow streams and risk exposures, as Exhibit 27-4 shows.

Similarly, zoning rules vary, with strong investor protection from strict zoning rules seen in some countries and regions, but paired with less growth, and less zoning protection but more growth in other countries like Poland, the U.S., or China. Diversification across these can also create a more stable and more predictable return.

There are also **theoretical underpinnings** to the idea of international diversification. Modern portfolio theory and the concept of the efficient frontier show how diversification can enable an investor to reach for a higher expected return with a given level of risk, or a lower risk with a given expected return. The efficient frontier we previously showed, represented the best possible portfolios in the home market. The question is how the position and shape of the efficient frontier would be affected by going international. Exhibit 27-5 provides a stylized graph of the national and global efficient frontiers. When we look at the global market, we find that the global efficient frontier dominates the national frontier, providing higher returns for a given risk. A bigger universe to choose from means greater diversity, so the correlations are lower, and the risk is reduced, hence

EXHIBIT 27-3

International Correlations of Selected Market Fundamentals across Economies, 1995–2022

Panel A: Correlation Matrix Annual GDP % Change

	Australia	Brazil	China	Euro Area	India	Indonesia	Japan	South Korea	U.K.
Brazil	0.19								
China	0.20	0.68							
Euro Area	0.40	0.19	0.09						
India	-0.17	0.04	0.39	0.03					
Indonesia	0.48	0.82	0.62	0.18	-0.04				
Japan	0.13	0.42	0.20	0.59	0.17	0.54			
South Korea	0.17	0.39	0.29	0.33	0.03	0.70	0.47		
U.K.	0.47	0.09	0.00	0.83	-0.04	0.34	0.70	0.32	
U.S.	0.57	0.08	-0.05	0.76	0.03	0.28	0.60	0.28	0.87

Panel B: Correlation Matrix Annual CPI % Change

	Australia	Brazil	China	Euro Area	India	Indonesia	Japan	South Korea	U.K.
Brazil	0.40								
China	0.41	0.81							
Euro Area	0.62	0.38	0.52						
India	-0.10	0.18	0.30	0.09					
Indonesia	-0.20	-0.05	-0.21	0.16	0.35				
Japan	-0.28	-0.10	0.11	-0.16	0.02	0.06			
South Korea	0.18	0.24	0.29	0.61	0.41	0.59	-0.03		
U.K.	0.28	0.12	0.50	0.62	0.42	-0.05	0.05	0.38	
U.S.	0.57	0.13	0.38	0.83	-0.27	-0.03	0.07	0.32	0.51

Panel C: Correlation Matrix Interest Rate

	Australia	Brazil	China	Euro Area	India	Indonesia	Japan	South Korea	U.K.
Brazil	0.17								
China	0.83	0.44							
Euro Area	0.68	0.83	0.84						
India	0.74	0.71	0.80	0.93					
Indonesia	0.27	0.91	0.43	0.79	0.68				
Japan	0.78	0.75	0.85	0.92	0.88	0.63			
South Korea	0.51	0.84	0.71	0.92	0.83	0.91	0.89		
U.K.	0.47	0.79	0.45	0.71	0.76	0.69	0.76	0.76	
U.S.	0.55	0.67	0.63	0.74	0.76	0.70	0.76	0.75	0.91

Source: Authors' calculations, based on data from World Bank and ECB.

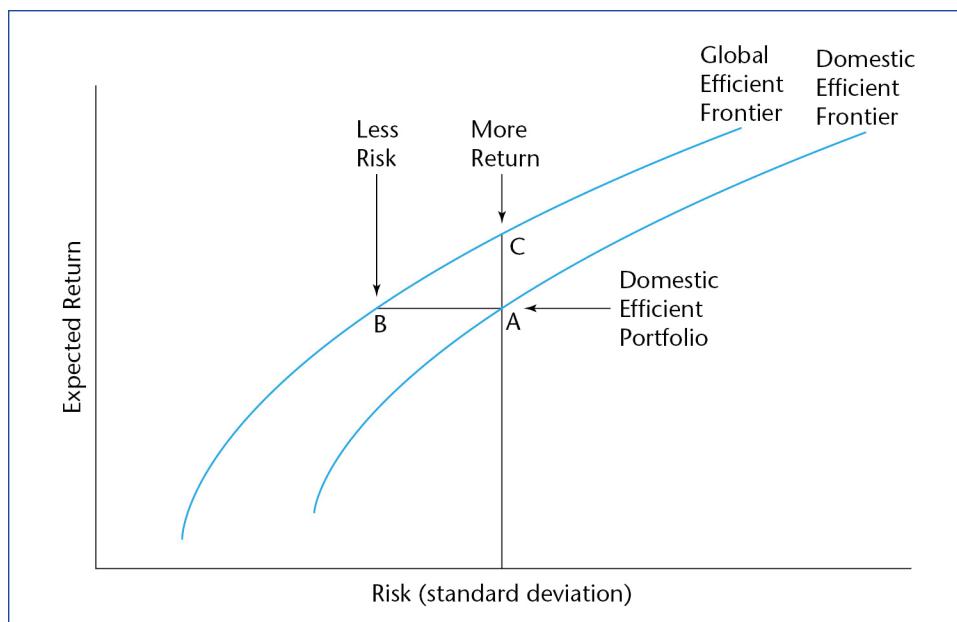
the lower risk of the minimum variance portfolio. As well, it will be possible for a U.S. investor to find a market with higher expected return—and higher associated risk—than would be possible in the American market alone, hence the higher expected returns at the right side of the frontier. Essentially, picking from a larger investment universe provides more diversity, more potential on the upside and more protection on the downside. As a result, the investor holding domestic portfolio A can reduce her risk by going to international portfolio B or improve her expected return by going to international portfolio C. She can also aim for a combination of these two goals by choosing a portfolio between points B and C on the efficient frontier.

EXHIBIT 27-4

Rental Contract Types in Selected Property Markets across the Globe

Country	Lease Length	Right to Renew	Rent Reviews	
			Period	Basis
Asia-Pacific				
Australia	3–10 years	Varies	Varies	Varies
China	2–5 years	Yes	Expiry	Market
Hong Kong	3–6 years	Yes	3 years	Market
India	3–9 years	Yes	3 years	Market
Japan	2–5 years	Yes	Expiry	Market
Korea	5 years	Yes	Annually	CPI
Singapore	3 years	Yes	Expiry	Market
Americas				
Argentina	3–5 years	Negotiable	Negotiable	Case-by-case
Brazil	5 years	Yes	End of term	CPI
Canada	5/10 years	Negotiable	5 year	Market
Mexico	3–5 years	Negotiable	Annual	CPI
United States	3/5/10 years	Negotiable	N/A	Fixed/CPI
Europe				
France	3/6/9 years	Yes	Expiry	Market
Germany	5–10 years	Yes	Annual	CPI
Netherlands	5–10 years	Yes	5-10 years	CPI
Spain	5–10 years	Yes	5 Years	Market
United Kingdom	5–10 years	Yes	5 years	Market (upward only)

Source: NAIdirect.

**EXHIBIT 27-5** Domestic versus Global Efficient Frontier.

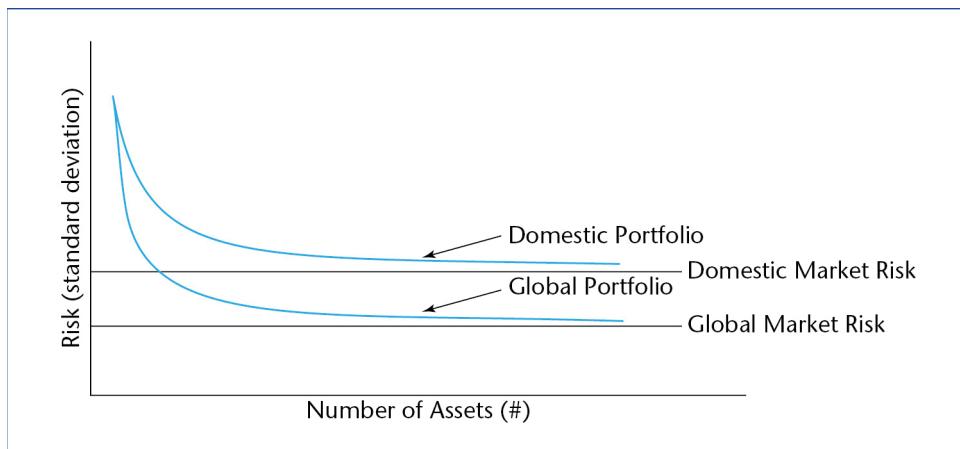


EXHIBIT 27-6 Domestic and Global Market Risk.

We already showed that the risk of an investment can be split into two parts: market risk and specific property risk. As more assets are added to the portfolio, the stronger the diversification, and the lower the risk, until the portfolio's risk equals market risk, and no further reduction is attainable. If the universe selected from is limited to the local market, the total risk can be reduced only to some extent. But if the market is widened to the global universe, this provides more diversity and more diversification potential. So, in a typical situation, the global market risk will be lower than the domestic market risk. A downward shift of market risk can be achieved through an increased universe of assets. This is not true at any point in time for any market, but holds generally, for a long-run investor. This means it is possible to diversify away a greater part of the total risk and run less risk for a given return. On these grounds, one could argue that all investors should be international.

Looking at Exhibit 27-6, one would be tempted to conclude that the main conclusion from the one-country CAPM could be easily extrapolated to a global setting and that all investors should hold the same global market portfolio. However, this is not necessarily the case. The decisive question here is whether capital markets are internationally integrated or segmented. If assets are priced in integrated markets, expected returns will be in accordance with the global systematic risk as depicted in the lower horizontal line in Exhibit 27-6. If, on the other hand, assets are priced in segmented markets, their returns will be in line with the systematic risk of their domestic market. Since this is generally higher, the expected return will be higher as well. This implies that an investor who can avoid the cause of this market segmentation will enjoy special benefits from international diversification.

So, the question is whether international asset markets are integrated or segmented. For stocks, the evidence points in the direction of increasing integration, and for the publicly listed real estate market, the same has been documented.⁹ But it is not likely that private real estate markets are as integrated as their public counterparts. Direct legal barriers to foreign real estate ownership, and more indirect informational barriers putting international investors at a disadvantage still exist, allowing sustained international price discrepancies. But as noted in the previous sections, markets are slowly becoming more transparent, and foreign ownership restrictions are gradually weakening, so price discrepancies will probably diminish over time.

⁹ Bardhan, Edelstein, and Tsang (2008) show that increasing international openness of the real estate markets leads to lower excess returns relative to the risk-free rate, even though this openness may also increase rents and asset prices.

From a theoretical asset pricing perspective, the difference between a country and the world is that a country has only one currency, while the world has more than one. The difference between the International Asset Pricing Model (the IAPM) and the traditional single-country CAPM therefore has to do with currency effects. The IAPM starts with the assumption that investors care about risk and return expressed *in their own currency*. Besides that, two more conditions are required: purchasing power parity holds and investors all have the same consumption basket, so inflation is measured the same way in all countries. All assumptions of the single-country CAPM are also applicable to the IAPM. If these assumptions are all valid, the outcome of the IAPM as a prescription of what investors should optimally do is very much in line with that of the single-country CAPM: investors should hold a portfolio consisting of risk-free bonds denominated in their own currency and the world market portfolio, optimally hedged against foreign currency risk. In other words, the separation theorem also holds for the IAPM.

However, the empirical problems associated with the single-country CAPM are also relevant for the IAPM, and probably even more so. Holding the global market portfolio is even more difficult than it is to hold the home-country market portfolio, especially when it involves private markets – as is still predominantly the case for real estate. As with the single-country CAPM, beta does not explain international return differences very well, and especially financial distress seems not to be captured well by the standard IAPM.¹⁰ But that is no reason to discard the model altogether. Global beta may not be the whole story when it comes to explaining asset returns on an international level, but it is a big part of the story, just like the beta of the single-country CAPM. Additional risk factors probably should be incorporated into the model.

The empirical evidence in the literature regarding international diversification is mostly supportive. The diversification argument critically depends on international correlations, on the question of how synchronized different markets behave. Generally, researchers find that international correlations are low: the correlations between real estate markets may be even lower than those between stock and bond markets. This would suggest a clear argument for international investment in real estate, stronger even than for international investment in stocks and bonds. This holds especially when looking at the global universe of property investment opportunities. Within Europe, on the other hand, correlations are quite high, nearly as high as across U.S. regions. But across Europe, North America and the Far East, correlations are low.¹¹ The evidence is not all supportive of international real estate diversification, though. Some authors find a dominant global factor driving real estate markets—both public and private—suggesting the international diversification effect is not so great. Others find increasing international integration between real estate markets.¹² However, looking at the body of empirical evidence regarding this matter, it seems safe to say that international diversification in real estate portfolios is a good idea.

Reviewing the case for international diversification, it is clear that the diversification effect depends on the correlations between international real estate markets. Thus, the question is how stable (thus, predictable) these correlations are. Unfortunately, they are not; they tend to move around, even over short horizons. That would probably be all right if the correlations merely fluctuated up and down, but if they move in a structural way against the investor, this is a problem. If correlations start low and then move in the direction of one, the whole diversification effect would dissipate.

¹⁰ Ling and Naranjo (2002) show that a global market beta is a significant driver of expected returns on listed property companies and Bond, Karolyi, and Sanders (2003) extend this research to a multi-factor framework.

¹¹ Such supporting evidence has been put forward among others by Eichholtz (1996), Conover, Friday and Sirmans (2002), Hoesli, Lekander and Witkiewicz (2003), De Wit (2010), Liow (2010), and Kroencke and Schindler (2012). The evidence shows that international correlations are relatively low, and especially so across continents and economic blocks. This results in significant benefits from international investment.

¹² Examples are Goetzman and Wachter (2001) and Ling and Naranjo (2002). Good literature reviews can be found in Sirmans and Worzala (2003a and b). More recent evidence for increasing integration can be found in Yunus (2012).

At first sight, one might expect globalization and capital market integration to mean that correlations get higher, and the diversification effect reduces. However, pure globalization is not actually what is happening; instead, there is a trend towards regionalization. Continental factors are important and are increasing in importance. Within the regional blocks, correlations are going up, but between regions, they are not. The most defined regional blocks are North America and Europe, while within Asia, the same effect can be seen but it is weaker: correlations are going up in a far weaker way than in Europe and the Americas.

The second question is whether the low correlation is there when it is needed. Diversification is essentially a form of risk management. So, when some markets are going down very rapidly, an investor wants to be protected by the others. Unfortunately, correlations tend to go up in times of economic crisis, which is when you most need them to be low. There is ample evidence showing this effect for stocks and bonds, but in real estate, the correlation increases in times of crisis appear to be somewhat less pronounced. So, the diversification effect in real estate is weaker in times of crisis, but still favorable compared to other asset classes. Finally, all these results depend on the investment horizon. If it is long enough, the investor can simply wait out the higher correlations in a crisis, making sure to have a sound international real estate position, with exposure to strong economies and demographics.

27.2.3 MANAGERIAL CONSIDERATIONS

The third possible rationale for going international consists of managerial considerations, like exporting management expertise, and servicing global customers. To start with the former, it seems obvious that the know-how in real estate investment management built up by investment managers in the mature property markets has a value and could be exported to emerging markets in Asia and Central and Eastern Europe, for example. With their skills at managing real estate portfolios and developing projects, one would think they could carry out projects that a local player might not be able to do with the same degree of success. Unfortunately, an advantage in know-how in the financial markets is inclined to evaporate rapidly, especially in view of international labor markets and the international market for higher education. Those working for a U.S. investor can also work for a local player and come up with the same projects. Without natural monopolies, advantages in know-how quickly disappear as knowledge is dispersed.

Another question is, how valuable is know-how in general in property investment? Publicly listed property investment companies are sometimes priced at a premium to net asset value—the value of the assets minus the debt, and thus the tangible assets in the company. But often, they are priced at a discount, suggesting that the market does not always perceive management as a value creator. This does not say much for the value of know-how, and if it does not have much value, there is not much to export, and the whole argument collapses. It would seem evident that there is important know-how in property development, but even that management premium is doubtful, as recent research suggests that only the land contracts in development have value, and not the management.¹³

The other managerial argument for becoming an international property investor is to follow one's customers. As corporations are becoming more and more global, so do the companies that service them. Property investors are part of the business services community. If corporations centralize their global purchasing of services, as they have generally done with accountancy and banking, for example, then it makes sense for the service providers to have the same global footprint as their clients.

¹³ The argument that management matters has the most impact for retail property where tenant mix is especially critical. Combining the right tenants will result in greater sales and higher rents, better utilization of parking spaces and a more valuable property.

To judge the strength of this argument for international property investment, two issues need consideration. The first is how international the property users really are, and the second is how central the corporate real estate decisions are made within these organizations. Concerning the former, it is likely that the sector in which the company operates plays a role. For example, the distribution and logistics industry is a global business. That means it makes sense that property companies providing space to that sector are global as well. The American firm ProLogis is a prime example of this strategy and seems to be doing well.

Retailing is not that international: hardly any of the great American retail chains have successfully set foot in Asia or Europe, and vice versa. However, retail seems to have become more international in the last decade. True global retail formulas are still rare, but European chains, traditionally confined to their own country, are branching out to other European countries. As a result, the property companies servicing them are becoming pan-European, and even global, as well.

The last two main institutional property types, offices and multi-family housing, are even less international than the retail sector. For offices, this may have something to do with a lack of global space occupancy decisions among office users. For housing, demand is simply not international. In short, the power of the 'follow your customers' argument depends on the property type but is generally not very strong.

27.3 OBSTACLES TO INTERNATIONAL PROPERTY INVESTMENT

However, there are also obstacles and risks attached to going international. It entails various costs that are difficult to recoup, and the foreign investor is likely to be at an informational disadvantage relative to the local competition. Partnerships or joint ventures may be one way to overcome this. Secondly, going international leads to additional risks, like currency risk and political risk, and since international investments are likely to be less liquid than local property investments, these risks could be harder to manage abroad than at home. Lastly, the international investor may face discriminatory regulations, like property ownership restrictions or adverse taxation. This section discusses these obstacles in depth.

We start with a close look at the costs of investing abroad, including transaction costs and a key element, information costs, where the remote investor is at a disadvantage. Liquidity is more of a problem for the international than for the local investor, and this is also true of political risks. Foreign investors are more vulnerable, due to their weak political clout. Other obstacles can include unfair laws and even corruption.

One obvious **cost** disadvantage, particularly for the U.S. investor, relates to property transaction costs abroad. In the U.S., property transfer tax, for example, scarcely exists, but in many other countries, it is quite high, particularly in countries like Spain (7 percent), the Netherlands (10.4 percent), India (5–7 percent), and Germany (5.5 percent). Exhibit 27-7 provides information regarding transaction costs for selected international property markets. Besides transfer taxes, investors are faced with agents' fees and legal fees. Agents' fees can also be high, while legal fees vary widely. This means the transaction costs of going international in real estate are both high and variable. Total costs for selling a property in countries like France, Brazil, and Spain, for example, will add up to about 10 percent. Investing abroad will require either a fantastic timing strategy or a long-term commitment, to recoup these extra costs.

However, the major issue here is **information costs**. These are extremely important in private real estate markets. There is overwhelming academic evidence showing that the public stock markets are rather efficient in an information sense, implying that the available and relevant information is reflected in asset prices. Having access to information then becomes meaningless, because the information is already incorporated in the prices, and it is not possible to outperform on that basis. This may sound disappointing, but the flip side of the argument is that it is not possible to underperform either, at least not due to a lack of information. In such markets, where all assets are fairly priced, it

EXHIBIT 27-7

Transaction Costs in Selected Private Real Estate Markets

Country	Transfer Tax	Agent's Fee	Legal Fees
Asia-Pacific			
Australia	5.50%	1.5–2%	0.5–1%
China	1.50%	1–1.5%	Each party bears their own legal costs
Hong Kong	0.01–3.75%	1%	Each party bears their own legal costs
India	5–7%	1–2% (seller) 0–0.5% (buyer)	Varies, each party bears their own legal costs
Japan	0%	3%	Varies
Korea	Varies(seller)/4.6% buyer	0.2–0.9%	Each party bears their own legal costs
Singapore	3%	1–2% (seller)	Vendor and purchaser pay their own legal fees
Americas			
Argentina	2.50%	1–4%	TDB
Brazil	2–3%	1–6%	2–3%
Canada	Varies by province	1.5–4%	N/A
Mexico	2%	5%	N/A
United States	Varies by state	1–6%	N/A
Europe			
France		Negotiation	Negotiation (usually 10%)
Germany	3.5–4.5%	1–5%	1.50%
Netherlands	10.50%	2%, payable by party represented	payable by each party
Spain	7%	6–15%	Negotiation
United Kingdom	4%	1%	Each party bears their own

Source: NAIdirect.

is safe for even completely uninformed investors to make big bets. They must only make sure they are properly diversified. This is the conceptual foundation of the thriving ETF markets in equity investment. It is also great news especially for international investors since they are likely to be the underinformed parties in any market.

Private real estate markets, on the other hand, do not fit this ideal model, and are not as efficient at incorporating new information into asset prices. The investment implication of this is that it increases the danger of doing a bad deal: paying too much, or receiving too little, and doing that in a consistent way. So, in the private real estate markets, having access to information is probably a key driver of performance. It is possible to beat the market consistently, and to be beaten by it consistently, depending on information or the lack of it. It is important to note here that being informed is a matter of degree. Players in domestic real estate markets probably have different degrees of market information, but it is likely that foreign investors are generally badly informed relative to most local investors. Logically, if there are both local and international investors in the market, the local investors will be the insiders, already holding the vital information, while the international investors are more likely to be outsiders without this information. If this is so, it should be reflected in performance.

There is relatively little research providing firm evidence, but the two existing studies do indeed show that this inefficiency issue is an important driver of performance in international real estate markets.¹⁴ What emerges is that international property companies as a group underperform mimicking

¹⁴ See Eichholtz, Koedijk, and Schweitzer (2001), and Eichholtz, Gugler, and Kok (2011).

portfolios of domestic investors by 2.7 percent a year, and this is consistent over time, applying even for sub-periods. Only during the last few years has this outperformance of the local investors decreased, which may be related to the increasing international transparency of property markets we discussed earlier. More recently, these adverse distance effects have also been investigated in the U.S. commercial real estate market, showing that far-away investors often pay too much when buying real estate, and face higher vacancy rates than their local competitors.¹⁵

This strongly suggests that information costs are indeed important performance drivers for international real estate investors. However, it should be noted that this reflects the average underperformance of the universe of international property companies. Within that universe, there were also companies that outperformed their benchmark. The results therefore suggest that information costs make it very difficult to perform well in foreign markets, but that it is not impossible.

The information costs can come in two varieties, both leading to underperformance. The first is that investors do not have the necessary information and therefore make mistakes. They buy lemons; they pay too much when they buy and get too little when they sell. Alternatively, they could try to solve the problem by buying information, for example from local brokers, or by establishing local offices and employing local people. However, that would translate the cost of the information disadvantage into the payment of fees and salaries, likewise eroding the return. In any event, these information costs imply that diversification, dubbed as the only free lunch in financial markets, is no longer free.

Along with problems of costs and liquidity, there are specific **risks** attached to international real estate investment. These can be grouped as political, economic and currency risks. **Political risk** looms larger in real estate than in any other asset class, with the possible exception of art, especially for foreign investors. This is because real estate and land are often considered part of the national heritage. Strong national emotions are then involved. For example, in the late 1980s, when Japanese investors were buying up companies in the U.S., the Japanese acquisition of the Rockefeller Center caused more furor than any other transaction at the time. These emotions can provide political support for continuous impediments to foreign real estate ownership. Switzerland is an example of a market that is hard to access for foreign real estate investors, and there can be similar ownership problems in Central and Eastern Europe. In China, ownership of land and buildings is still a serious challenge, disfavoring foreign investors. Although there may be ways around these types of problems, that is, long-term land leases, they require creativity and partnerships. There may also be impediments to repatriation of profits from foreign direct investments, including direct real estate, again requiring creativity, or simply not investing.¹⁶ Although these issues are gradually becoming less important, especially within trading blocs, this process will take a long time. In the case of the possible problem of double taxation, this is disappearing more rapidly as more and more tax treaties are concluded, generally aiming to eliminate double taxation.

Besides political risk at a national level, real estate markets are also vulnerable to local politics. In all real estate markets, local government will be extensively involved in matters such as zoning and tenant protection. This is far-reaching and local government is not going to go away. Foreign real estate investors are far more vulnerable because of their weak political clout, especially without local partners. They do not represent any voters, seldom contribute to election campaigns, and are not big employers.

Here the information asymmetry issue and political risk resurface. One of the big drivers of real estate performance is land use, and the remote investor is not aware of what is cooking politically, either at the level of central government, or at City Hall. The home market investors, on the other

¹⁵ See Ling (2018) and Eichholtz, Holtermans, and Yönder (2016).

¹⁶ For example, some U.S. investors have teamed up with Chinese investors such that instead of taking profits from China, they exchanged these with Chinese partners who had profits in the U.S., as a way to bypass the difficulties imposed by the Chinese government.

hand, will have a much better chance of knowing what is currently being discussed, or promoted, in the corridors of power and as well as tentative rules and regulations forthcoming. Land could be changed from industrial zoning to commercial use or from pasture to housing, providing highly profitable opportunities to some, and possible risks to others. To capitalize on local knowledge, the investor needs to be a household name at City Hall, and this is all but impossible for the remote investor, who may not even speak the local language. There may also be specific unfair laws against foreign ownership, or laws may be changed against foreign ownership without proper notice to the foreign investor.

Economic risk is partly diversified away by foreign investment, but foreign economies may in themselves be more volatile than the domestic market, especially if the domestic market is mature and the targeted markets are emerging. In that case, the higher returns expected in these markets will be accompanied by higher risks, and the investor will move further to the right side of the efficient frontier.

Investors will also be exposed to **currency risk** if they go abroad, and this could be far bigger than the property risk: volatility in the Eurozone and American property markets, for example, has historically been lower than the volatility of the euro–dollar exchange rate. In the long run, this is likely to change, due to the general tendency for countries to merge into blocks. The European trade bloc has gone all the way by merging its currencies into the Euro. Although the 2011/2012 travails of the Euro countries do not exactly provide a shining example of how to run a common currency, it is still likely that other regions will also move to link their currencies together. This does not mean that currency risk between blocks will disappear, but that it will decrease within those blocks, provided these currency links are properly managed.

TEXT BOX 27.2 THE DOS AND DON'TS OF CURRENCY RISK MANAGEMENT

A property investor who crosses the border is faced with currency risk, no matter whether her holdings are public or private. The dollar value of foreign property holdings may go down independently of real estate market movements in the host country, just because the local currency weakened against the dollar. International property investors have to deal with this risk. Over time, different techniques and policies have been suggested to accomplish this.

In the 1970s, when the global regime of fixed currencies was abolished, economists and econometricians aimed to predict currency movements. These predictions could then supposedly be used by corporations and investors as a basis for a currency policy. The idea was to hedge only when the foreign currencies were predicted to fall. However, abundant research since then has shown that currencies mostly follow a random walk, and that currency movements are hard, if not impossible, to predict.

In these circumstances, property investors who decide to run currency risk when they expect exchange rates to move in their favor do not serve their clients well, since they take on speculative positions in a market they probably do not know much about. Besides that, these currency bets take valuable management time, which could better be spent making property decisions. If currency movements are indeed unpredictable, then international property investors will have to hedge not at all or always and should not make this decision dependent on expectations regarding currency movements¹⁷.

¹⁷ In that respect, currency hedging is just like fire insurance. Since it is impossible to predict if and when your house will burn down, you do not take out insurance on the basis of your expectations regarding this. People rarely regret having paid their insurance premium when their house failed to go down in flames.

There are good arguments for not hedging at all. First, there are costs involved in hedging, as in any insurance product. The costs equal the forward premium, which varies across currencies and over time. Using options for currency hedging is even more expensive, since they cover the investor on the downside, while maintaining her exposure to the upside. Second, the international efficient frontier is not very strongly affected by the question of whether returns are calculated in the home currency (representing the unhedged situation) or in local currency (representing the hedged situation), especially in the long run. In a broad international portfolio, covering many countries and currencies, the currency risk is partly diversified away. Last, it is very hard to fully hedge the currency risk of foreign investment returns in practice, since these returns are not fully predictable, which makes it unclear how much of a hedge one needs to take out at any given time.

On the other hand, the arguments for always hedging are also rather convincing, especially for real estate. First, currency risk does increase investment risk, especially for portfolios covering a few countries, and even more so if the countries involved are emerging markets, which often have very volatile exchange rates. Second, institutional property investors are usually looking for a relatively modest risk-return profile, while currency returns can fluctuate rather strongly. For example, the volatility of the euro/dollar exchange rate has historically been much higher than the volatility of the property markets on either side of the Atlantic, so the targeted risk-return profile may well be thrown on its head because of the currency movements. Third, the development of deep and liquid currency derivatives markets in the last two decades has enabled investors to hedge rather cheaply and quickly, although this is less so for emerging market currencies.

An inexpensive and easy way to hedge a lot of currency risk is to acquire “natural” hedges through local leverage. For example, a property investor with a 40 percent overall leverage ratio could do all her borrowing at home, but from a currency hedging point of view, it would be far cleverer to borrow locally and apply that same leverage ratio to each individual country in the portfolio. That way, the remaining currency exposure would only be as big as the local equity position.

To hedge this equity currency exposure, an often-used technique is currency overlay, in which the portfolio is regarded in full and the currency diversification effects are included in the analysis. The investor then decides what percentage of the remaining exposure—after leverage and diversification—she wants to hedge. In effect, an overall risk-return trade-off is made for the currency exposure.

From a cost–benefit point of view, it is wise to concentrate one’s hedging on the liquid currencies where possible. For example, the liquidity in the U.S. dollar / Swedish kroner market is higher than in the euro / Swedish kroner market, translating into much lower bid-ask spreads. For a Swedish investor who wants to hedge the currency risk on her Euro-area real estate portfolio, it is cheaper to buy dollars for kroner, and subsequently buy euros for these dollars, than to buy euros directly. Likewise, a Euro-area investor who wants to buy properties in New Zealand is also better off by going through the U.S. dollar than buying New Zealand dollars directly. However, these bid-ask spreads change all the time, so the optimal route needs to be determined at the moment the hedge takes place.¹⁸

¹⁸ The three leading textbooks on international finance and investments are Sercu (2009), Solnik and McLeavy (2008), and Levi (2009). All three provide detailed insights into currency hedging. Sercu is the more analytical of the three. Solnik and McLeavy focus mostly on investments, and Levi is more concerned with international finance.

To end this section, we need to deal with **liquidity**. Always a problem in property markets but getting more important in far-away markets because the investor will be last in line with the local broker and will be less successful at monitoring her performance. Especially when pursuing an active trading strategy, this may be problematic. It will be difficult to get into a distant market in times of a boom, as the local people will have seen all the good deals, and the international investor is likely to be left with the ones that everyone else has explored and decided are not worth the risks. The benefits of arbitraging between markets can be a challenge when selling takes time. For example, selling offices in New York when the market is at the top of the cycle to move to London where it is at the bottom may take some time, especially if other investors have similar expectations about the New York market. It will be equally difficult to get out of a market when it turns down. The low liquidity of real estate makes the political and economic risks harder to escape. Even if a foreign investor sees economic hardship for the host country on the horizon, she may not be able to liquidate her holdings in time to avoid it.

Some of these risks and obstacles are diminishing in importance through increasing transparency, while others are inherent to international real estate investment. In that case, the international investor either must develop a strategy dealing with these issues, or retreat to the domestic market. The next section will introduce you to some such international investment strategies.

27.4 DEVELOPING AND IMPLEMENTING INTERNATIONAL REAL ESTATE STRATEGIES

Given all the obstacles and risks involved in international investment, how can an investor achieve a satisfactory performance? A successful strategy for international real estate will take advantage of the benefits, while avoiding as far as possible the obstacles and risks. The way to do this is to determine the key obstacles and risks that are more cumbersome for international than for domestic players and draw up a strategy to deal with them. They will include information costs, monitoring costs, currency risks, and political risks plus the extra problem of liquidity. The first important resource that can overcome many of these problems is the global property share market. REITs and other listed property companies are mostly local specialists, focusing on their own country. Can they overcome the obstacles and risks we have described? What role can REIT investment play in an international strategy? After having answered these questions, we can draw up our strategy, including the question of allocation, and introduce the ‘home market’ concept.

27.4.1 PUBLIC REAL ESTATE INVESTMENT

To establish how the global REIT market can help the international real estate investor, we will review the obstacles to international investment we have described and see how REITs measure up. On the costs side, transaction costs for real estate shares (and for shares in general) are low and getting lower. And since the property share market is rather efficient, information costs are low or non-existent. Anybody could make a well-informed property share deal in, say, Hong Kong by simply buying shares on the local stock exchange. But if they wanted to buy a private property in Hong Kong, this would be a totally different matter. Listed property companies enable uninformed investors to make well-informed foreign property bets. To make that situation even better, property companies are generally improving their reporting quality, partly forced in that direction by industry associations, but also by accounting rules.

As well, monitoring a foreign real estate portfolio is easy, as investors can free ride on the local co-shareholders and public information; there will be no need to rush over to Hong Kong to look at your property. Finally, impediments to foreign ownership scarcely arise in the REIT market, as they are simply shares.

Turning to liquidity, this is far better than for direct real estate holdings and comparable to shares in companies of similar size. Besides that, there is a non-linear relationship between size and liquidity, and since listed property companies are generally growing in market capitalization, their liquidity will go up even more. Discrimination against the foreigner in entering or exiting the market is not likely for listed real estate.

While the economic risk of public property investments is as large as for direct investment, greater liquidity means that investors are better able to defend themselves by withdrawing from the market. The same goes for political risk, with the extra advantage that the manager of the local property company will have the contacts and political clout that the remote direct property investor lacks. Finally, problems relating to regulation will largely disappear, as foreign ownership constraints and barriers to repatriation of foreign-earned returns are usually small or non-existent for stocks.

In short, the obstacles to international property investment discussed in the preceding section are either not relevant or less important for listed property companies, so they look like an ideal investment channel for international property investors. The proliferation of listed property companies has also spurred the development of the market for (international) mutual funds investing in their shares, as well as Exchange Traded Funds tracking listed REIT indices. These funds mainly cater to private investors and smaller institutions, and mostly invest in locally operating property companies, thus offering the benefit of one-stop shopping for investors who want to build up exposure to the global property share market.

27.4.2 DETERMINING COUNTRY ALLOCATION

For an investor who has decided that the global property share market is indeed the best way to build up international property exposure, a practical question is how to determine optimal portfolio weights. There are several ways to accomplish that. The first is to use modern portfolio theory to find optimally diversified international portfolios, the second is to track international index weights, and the third is to use weights based on GDP.

Using modern portfolio theory, a Markowitz framework or optimizer can be used to establish optimally diversified portfolios in the standard risk/return trade-off. However, the theoretical basis for using that approach within the real estate portfolio is weak, and this approach also encounters a few practical problems, like data availability. Studies using this approach have mostly been based on historical time series, and the resulting optimal portfolios were usually very period-specific, making them not very useful in practice. They also tend to deliver what are called corner solutions, which are often very time dependent. For example, the model's optimal outcome for high-risk / high-return investors is usually to allocate 100 percent to the asset or country that happened to have the highest return and risk during the sample period. This is hardly a solution for international diversification.

The logical alternative is to track the composition of the global market using the market weight of the global index. In international equity investment, tracking indexes like the global MSCI Index is a widely accepted and frequently used approach. For property companies, such indices are also available: FTSE EPRA/NAREIT, GPR and the S&P Global REIT Index are leading examples. But basing the portfolio weights on the index weights is liable to the coincidence that market weights are high in countries that happen to have a well-developed property share market. In Europe, for example, the U.K. pulls more than its weight in listed property shares, while Germany is underrepresented. In the Asia-Pacific region, the same is true for Australia, which has a dominant property share market, and a relatively small economy. Tracking a global index without making any adjustments for this will probably create unwanted specific risk.

The wisest solution, then, could be to use a combination of market index weights with GDP weights. Investors commonly possess an intuitive map of what the global market should look like, which is partly based on what the capital market looks like, but probably also partly based on what economies look like, for which GDP weighting can be used. Indices have been created in line with

this idea, allowing international investors to measure their performance when choosing their international allocation in this way.

27.4.3 THE “HOME MARKET” CONCEPT

We have seen that all the obstacles to international real estate investment are much reduced or disappear in the case of investing in property companies. However, well-versed real estate investors with experience may have an information advantage in their home market, which could enable them to outperform in direct real estate.

Besides that, there is some concern regarding the diversification effect of listed property shares in the broader mixed-asset portfolio. The early empirical evidence concerning the performance of listed property companies showed that the correlations between property shares and common stock tended to be high, while those between property share returns and direct real estate returns were found to be low. In other words, property shares seemed to behave more like stocks than like property, making them less attractive as a portfolio diversifier. Even though more recent empirical research results indicate that the correlations between property shares and common stock have decreased and seem to stabilize at lower levels, there exists no clear theoretical argument why this should be the case, so it may be just a temporary statistical fluke.

Moreover, an international portfolio that consists only of listed property shares may be optimized in terms of country allocation, but very unbalanced in property sector exposure. The degree to which property sectors are securitized varies strongly across countries, so sector compositions differ a lot across national property share markets. For example, the U.S. market provides an investor access to a variety of alternative property types, like self-storage, data centers, casino's, prisons, and healthcare facilities, while these are hardly available in listed markets outside of the U.S.

Both these issues provide rationales to keep at least part of the real estate portfolio invested in private property.

Given that direct property markets are less efficient, investors should concentrate their *direct property investments* on market(s) in which they have an information advantage: their home market. This does not necessarily need to be defined in a regional way. It could also be that they have expertise, information, and valuable networks in a certain property type. So, the home market for an investor could be logistics properties in Northern Europe, strip malls in the U.S., or multi-tenant offices in Southeast Asia. It could also be defined in an even narrower sense, for example, offices in Brussels, or conference centers in Singapore. No matter how it is defined, the key criterion must always be the potential for consistent outperformance, driven by consistently better access to private information than the competition.

Deciding which markets can be considered true home markets in that sense is a difficult process, in which honesty is required, and ego needs to remain in check. After all, we all like to regard ourselves as potential outperformers. However, there is a simple line of thinking to help this process. The consistent outperformance you strive for will have to be accomplished at the expense of the consistent underperformers: the underinformed investors. So, a home market can be defined as a market in which you can point out the consistent underperformers. If you cannot, you are not an insider, and you'd better not become a direct investor in that market. In that respect, real estate investment is just like playing poker: at any poker table, there is a less experienced or naïve player, and if you cannot point him out, you know who it must be.

An important counter argument against this idea of concentrating the investments on the home market is of course diversification. The fact that information costs exist does not nullify the usefulness of spreading your risk internationally. It only means that there is a trade-off between the advantages of international diversification and the information costs this entails. The question, then, is what position an investor should choose on this trade-off.

The answer to that question depends on the nature of the investor. Let's first consider the intermediary investor, like a listed or unlisted property company. For such an investor, diversification does not add value since the shareholders on whose behalf the organization is making investments can most likely diversify their holdings. That implies that the soundest investment strategy for an intermediary investor is full concentration on the home market, thereby enabling outperformance of the competition. Most REITs seem to do that, as they tend to specialize in property types, and do not invest much abroad.

Nonetheless, where the diversification argument may not hold true for a property company, it will still be important for end-investors like pension funds, sovereign wealth funds, or family trusts. These end-investors do not have shareholders who can diversify their exposure to them, so they must diversify themselves. For such investors, it may still pay to determine a home market, or home markets, for the reasons given above. If they do, they should try to manage information costs with a true local presence in the home market(s) and build critical local mass to get access to the best deals and to the information flow required to outperform. This means an effective strategy in private real estate markets rules out covering too many markets. Unlike the intermediary investors, they should not stop there, though. The diversification requirement makes it wise to build up an international portfolio of listed and locally operating property companies to create exposure outside the home markets. Alternatively, the end-investors may decide that diversification is more important than the potential for outperformance and hold all real estate exposure through listed property companies.

It is now time to broaden the view a bit and look at other asset markets besides real estate markets. The analysis in this section and the previous ones was fully focused on diversification issues within the real estate portfolio. But as we argued before, portfolio considerations should not be limited to real estate holdings alone and should instead be taking the overall wealth portfolio of the investor into account. Following the same line of reasoning provides slightly different investment advice than the previous paragraph, since it is not clear whether a private real estate portfolio in the home market should be complemented with listed property shares or plain common stocks and bonds to reach an optimum on the trade-off between outperformance and diversification. To find the optimal portfolio besides the private real estate exposure in the home market(s), therefore, one must take all investment possibilities into account.

27.5 CHAPTER SUMMARY

This chapter provides an introduction to the global real estate capital market. Our goal was to introduce you to the rationale for international investment in real estate, but also to the obstacles for doing that. We then gave some suggestions for international real estate investment strategies that make sense given the nature of real estate assets and markets. These include using publicly traded REITs in foreign countries or finding local partners.

The reasons for investing in real estate outside one's own country are threefold. First, there may be return opportunities abroad (structural or opportunistic) that are unattainable at home. Second, since global real estate markets move in non-synchronous ways, there are diversification effects related to international investment. The third rationale for going international may be to export know-how in developing and/or managing real estate assets. However, going international can be cumbersome, due to low international liquidity, additional transaction and information costs, and new risks, like currency risk, political risk, and economic risk.

These issues suggest that international real estate investment is a trade-off aimed at maximizing its potential gains, while managing the disadvantages as much as possible. Different types of investors will do that in different ways. End-investors, like sovereign wealth funds and pension funds, should aim for diversification, acquiring direct property exposure in a few markets they know well, and using listed and unlisted property companies besides stocks and bonds in other markets.

Intermediate investors like property companies, for whom diversification does not add value, should focus on and aim for outperformance in their home market.

KEY TERMS

- International diversification
- (International) correlation
- Global market portfolio
- Home market
- Market timing
- Demographics
- Information costs
- Informational inefficiencies
- Real Estate Investment Trust (REIT)
- Property Trust (PT)
- Currency risk
- Currency risk hedging
- Economic risk
- Political risk

STUDY QUESTIONS

Conceptual Questions

- 27.1. Why do real estate investors invest internationally?
- 27.2. How should an international property investor determine her country allocation?
- 27.3. In what respect is international property investment fundamentally different from international investment in listed stocks?
- 27.4. How can the global property share market help an international property investor?
- 27.5. Currency risk is an added risk factor in an international property portfolio. How should an investor deal with this risk?
- 27.6. Discuss the reluctance of institutional investors to invest internationally in real estate, considering risk exposure and the nature of informational (in)efficiency in the real estate asset market.

REFERENCES AND ADDITIONAL READING

Bardhan, A., R. Edelstein and D. Tsang. 2008. Global Financial Integration and Real Estate Security Returns, *Real Estate Economics*, 285–312.

Bond, S., A. Karolyi and A. Sanders. 2003. International Real Estate Returns: A Multifactor, Multi-Country Approach, *Real Estate Economics*, 481–500.

Brounen, D., P. Eichholtz and D. Ling. 2007. Trading Intensity and Real Estate Performance, *Journal of Real Estate Finance and Economics*, 449–474.

Conover, C., H. Friday and G. Sirmans. 2002. Diversification Benefits from Foreign Real Estate Investments, *Journal of Real Estate Portfolio Management*, 17–25.

De Wit, I. 2010. International Diversification Strategies for Direct Real Estate, *Journal of Real Estate Finance and Investments*, 433–457.

Eichholtz, P. 1996. Does International Diversification Work Better for Real Estate than for Stocks and Bonds?, *Financial Analysts Journal*, 56–62.

Eichholtz, P., N. Gugler and N. Kok. 2011. Property Market Transparency and the Costs of International Investment, *Journal of Real Estate Finance and Economics*, 152–173.

Eichholtz, P., R. Holtermans and E. Yönder. 2016. The Economic Effects of Owner Distance and Local Property Management in U.S. Office Markets, *Journal of Economic Geography*, 781–803.

Eichholtz, P., R. Huisman, K. Koedijk and L. Schuin. 1998. Continental Factors in International Real Estate Returns, *Real Estate Economics*, 493–509.

Eichholtz, P., K. Koedijk and M. Schweitzer. 2001. Testing International Real Estate Investment Strategies, *Journal of International Money and Finance*.

Eichholtz, P., M. Korevaar and T. Lindenthal. 2023. Long-term Growth and Urban Housing Rents. Cambridge University Working Paper.

Eichholtz, P. and T. Lindenthal. 2014. Demographics, Human Capital, and the Demand for Housing, *Journal of Housing Economics*.

Francke, M. and M. Korevaar. 2022. Baby Booms and Asset Booms: Demographic Change and the Housing Market. SSRN.

Gallo, J. and Y. Zhang. 2010. Global Property Market Diversification, *Journal of Real Estate Finance and Economics*, 458–485.

Gerlach, R., P. Wilson and R. Zurbruegg. 2006. Structural Breaks and Diversification: The Impact of the 1997 Asian Financial Crisis on the Integration of Asia-Pacific Real Estate Markets, *Journal of International Money and Finance*, 974–991.

Goetzmann, W. and S. Wachter. 2001. The Global Real Estate Crash: Evidence From an International Database, in: S. Brown and C. Liu, *A Global Perspective on Real Estate Cycles*, Kluwer Academic Publishers, 5–23.

Glaeser, E. and J. Gyourko. 2005. Urban decline and durable housing". *Journal of Political Economy*.

Gupta, A., V. Mittal and S. van Nieuwerburgh. 2024. Work from Home and the Office Real Estate Apocalypse. SSRN.

Hoesli, M., J. Lekander and W. Witkiewicz. 2004. International Evidence on Real Estate as a Portfolio Diversifier, *Journal of Real Estate Research*, 161–206.

Kroencke, T. and F. Schindler. 2012. International Diversification with Securitized Real Estate and the Veiling Glare from Currency Risk, *Journal of International Money and Finance*.

Lapier, T. 1998. Cross Border Direct Investments, Chapter 2 in: *Competition, Growth Strategies and the Globalization of Services; Real Estate Advisory Services in Japan, Europe and the U.S.*, Routledge, London.

Levi, M. 2009. *International Finance*, 5th edition, Routledge.

Lieser, K., and A. Groh. 2011. The Attractiveness of 66 Countries for Institutional Real Estate Investments, *Journal of Real Estate Portfolio Management*, 191–212.

Ling, D. and A. Naranjo. 2002. Commercial Real Estate Performance: A Cross Country Analysis, *Journal of Real Estate Finance and Economics*, 119–142.

Ling, D., A. Naranjo and M. Petrova. 2018. Search Costs, Behavioral Biases, and Information Intermediary Effects, *Journal of Real Estate Finance and Economics*, 114–151.

Liow, K.H. 2010. Integration between Securitized Real Estate and Stock Markets: A Global Perspective, *Journal of Real Estate Portfolio Management*, 279–266.

Marcato, G. and T. Key 2005. Direct Investment in Real Estate, *The Journal of Portfolio Management*, 55–69.

Sercu, P. 2009. *International Finance, Theory in Practice*, Princeton University Press.

Sirmans, C.F. and E. Worzala. 2003a. International Direct Investment in Real Estate: A Review of the Literature, *Urban Studies*, 1081–1114.

Sirmans, C.F. and E. Worzala. 2003b. Investing in International Real Estate Stocks: A Review of the Literature, *Urban Studies*, 1115–1149.

Solnik, B. and C. McLeavey. 2008. *Global Investments*, 6th edition, Pearson – Prentice Hall.

Yunus, N. 2012. Increasing Convergence between U.S. and International Securitized Property Markets: Evidence Based on Cointegration Tests, *Real Estate Economics*, 383–412.

Zhou, J. 2011. Downside Risk Spillover among Global Real Estate Securities Markets, *Journal of Real Estate Portfolio Management*, 255–270.

28 Real Estate, Sustainability, and Climate Risk

CHAPTER OUTLINE

- 28.1 The Setting
 - 28.1.1 Real Estate and the Environment
 - 28.1.2 Regulatory Issues: Government Policies for a Green Real Estate Sector
- 28.2 Financial Considerations for Owners and Developers
 - 28.2.1 Does Energy Efficiency Add Financial Value?
 - 28.2.2 The Evidence: Cash Flows, Discount Rates, Values
 - 28.2.3 The Timing of Energy Efficiency Investments
- 28.3 Physical Climate Risk and Real Estate
 - 28.3.1 Location, Location, Location
 - 28.3.2 Physical Climate Risk: The Big Six
 - 28.3.3 Effects, Risks, and Insurance
- 28.4 Chapter Summary

LEARNING OBJECTIVES

After reading this chapter, you should understand:

- The role of real estate in climate change.
- The financial considerations of energy efficiency in buildings.
- Climate risk for real estate, both transitional and physical.
- The consequences of climate risk for the functioning of real estate markets.

Real estate has a big impact on climate change, and therefore on climate risks, not to mention environmental risks in general. Most of us are concerned about pollution and the impact of pollution on our health. Just imagine spending some time in New Delhi, Mexico City, or the Los Angeles of the 1960s, and it is easy to understand that urban environments are not always healthy. Aside from pollution, we are concerned about weather-related property damage from wind, flood, and fires. These are part of the story of climate change. It does not matter here how much humans have contributed to climate change. It does matter that climate risks have increased and that real estate is part of the story.

The construction of a building directly produces a lot of CO₂, and once the building is in use, it typically needs gas and electricity, often produced by non-renewable sources, thereby also producing CO₂. The good news is that the industry can have a very substantial impact on reducing climate change and has begun working towards this goal.

This chapter will start by looking at industry initiatives aimed at improving environmental performance. We will then discuss regulatory issues, as governments become more involved, and will explore the transition risks for buildings that do not adhere to the changing standards. We will subsequently focus on what owners and developers can do to improve the environmental performance of buildings, focusing on energy performance. We use the net present value framework to work through

the financial ramifications of investments into buildings' energy efficiency, based on the available academic evidence, and show that green buildings are—largely—simply good business.

However, even if all commercial buildings were to achieve the environmental performance necessary to reach the Paris climate goals, the real estate sector would still face physical climate risk issues. The third part of the chapter, therefore, considers physical climate risk and its consequences for the industry worldwide. The 2015 Paris accord aims to restrict global warming to 2 degrees Celsius relative to the historic average, but while we are merely 1 degree warmer as of 2024, the world is already seeing the effects of more flooding, droughts, and wildfires. There is little doubt that this is going to get worse. Precisely because real estate is immovable, it is exposed to these climate events, presenting a set of new systemic risks to the industry.¹

28.1 THE SETTING

28.1.1 REAL ESTATE AND THE ENVIRONMENT

Buildings contribute to CO₂ emissions, both during and after construction. Using current standard technologies, cement and concrete production emit lots of CO₂. And when buildings are in use, they emit CO₂ as well, either directly by burning gas for heating, or indirectly, by consuming electricity. Overall, the real estate sector accounts for an estimated 30–40 percent of global energy demand and greenhouse gas emissions. In that sense, real estate is a polluting industry. It is also an essential industry. Humanity cannot survive without real estate, but the industry needs to change to become part of the solution to climate change. The strong potential offered by the built environment for greenhouse gas abatement is confirmed in IPCC reports from 2007 onwards.

In terms of environmental externalities, CO₂ emission by existing buildings is predominantly “scope-3 emission.” While buildings are *users* of electricity, the “scope-1 emission” resulting from electricity *production* actually occurs at the power station. If investors look at the real estate in their portfolio, they might therefore reason that this is not a concern for them.

However, energy costs affect both investors and tenants. Energy consumption in commercial real estate has been estimated to account for 30 percent of operating expenses and 7 percent of total occupancy costs. As energy costs go up—for example, from wars that disrupt or damage fossil fuel supply chains—the financial importance of this increases.²

To help developers, investors, and tenants make informed decisions, the market has created information tools measuring buildings' energy performance—often preceding government initiatives in that regard. Examples are BREEAM, LEED, and Australia's Green Star. These market-based measures tend to take a holistic view on buildings' environmental performance, also looking at things like access for pedestrians, closeness of public transportation, the availability of bike shelters and green roofs.³

These measurement systems are all at the building level. At the portfolio level, the market has created GRESB, aiming to measure the energy performance of (listed and unlisted) property companies and funds, on behalf of their shareholders.

¹ See www.unepfi.org/wordpress/wp-content/uploads/2024/04/Climate-Risk-Landscape-2024.pdf among many relevant climate assessments.

² Examples are the wars in the Middle East in the 1970s, and the Russia–Ukraine war from 2022 onward.

³ Unfortunately, this approach does seem to blur the relationship between these labels' rating levels and actual energy use. The correlation between LEED score and energy use in U.S. office buildings, for example, is very low.

TEXT BOX 28.1 ENVIRONMENTAL PERFORMANCE TOOLKIT

Market-initiated environmental performance tools, building level

BREEAM (Building Research Establishment Environmental Assessment Method) in the UK was the first environmental performance yardstick for buildings, established in 1990. It looks very broadly at buildings' environmental performance, going far beyond energy consumption alone. Buildings get rated in five categories from Pass to Outstanding. Also used in continental European commercial property markets.

LEED (Leadership in Energy and Environmental Design) followed in the U.S. in 1998. Its design was more or less taken over from BREEAM and has since then diverged but remains broad in scope. Buildings get points that are compiled into four certification levels ranging from Certified to Platinum. Also used in Canada, Asia, and Europe.

Green Star of Australia looks at the operational performance of buildings across nine environmental impact categories. Buildings can obtain between one and six stars, with six representing "world leadership" in environmental performance.

Market-initiated, portfolio level

GRESB (the Global Real Estate Sustainability Benchmark) measures the energy performance of (listed and unlisted) property companies and funds, on behalf of their shareholders. It not only looks at the actual aggregated energy use and carbon emissions of these companies, but also at their sustainability practices.

GREEN (the Global Real Estate Engagement Network) provides a platform and information toolkit for asset owners to engage with the property companies in which they invest, and for these property companies to engage with tenants, all with the aim of improving energy performance.

Government-initiated, building level

Energy Star in the U.S. is the oldest government-initiated energy assessment tool for buildings. It only looks at a building's actual energy use. It provides a relative score, and buildings only get an Energy Star certification when their energy performance is in the top-25 percent of similar buildings.

Europe's **EPC** (Energy Performance Certificate) looks at the theoretical energy performance of a building, based on its design. Buildings get points, which get converted into levels that initially ranged from G to A, now extended to A++++. More used for residential than for commercial property.

The **Green Building Evaluation Label** of China looks at six categories of environmental performance, awarding a maximum of three stars for overall performance. The label does not have a very wide market coverage yet.

CRREM (the Carbon Risk Real Estate Monitor) measures the risk of falling behind the energy efficiency goals from the watershed 2015 Paris climate agreement, and the resulting transition risks for individual assets and for real estate portfolios. It was originally an EU initiative and was later adopted and radically expanded by a global consortium of leading asset owners. It is used at the individual building as well as on the portfolio level.

There is broadening market awareness that something must be done about emissions and energy use, and that this is in the real estate industry's own interest. For example, *Corporations* are increasingly choosing sustainable buildings, either as tenants or as owners. This is often part of their social responsibility policy, but there can also be pressure from their workforce, or simply the desire to boost the company's labor market attractiveness. For example, Deloitte chose Amsterdam's extremely energy-efficient Edge building because it reasoned that in competing for top-talent with its rivals, a sustainable building would help attract the right people.⁴ In general, firms with high human capital intensity are most likely to use green buildings.

Furthermore, while the rent or the financing costs of a building are relatively stable, energy prices are very volatile. An energy-efficient building implies lower energy costs, meaning a more predictable cash flow as it reduces exposure to energy-price volatility.

Institutional investors, the ultimate asset owners, are also focusing on the energy performance of their real estate portfolios. For them, this is often a cornerstone of their ESG (Environmental, Social, Governance) policy.⁵ They use portfolio-level information tools like GRESB and CRREM to assess the sustainability performance of the property companies—mostly unlisted property funds and listed REITs—they invest in, and based on that information, they engage with these property companies to improve environmental performance.

Building owners are increasingly using these information tools to measure the energy efficiency of building design and construction. They also use these tools as yardsticks to improve their properties. In the U.S. office market, for example, the use of LEED certification has risen steadily from just 5 percent in 2005 to 40 percent in 2014. LEED market penetration is also strong in the U.S. retail sector. Energy performance certification is becoming a minimum sanity check in institution-grade real estate markets in many developed economies. In the rest of the world, however, this transparency is not available, as Exhibit 28-1 clearly illustrates. Especially in Asia, Latin America, and Africa, energy efficiency certification remains a challenge.

So, it seems that the sector is already taking steps to reduce its carbon footprint. While that may be due to a moral imperative of ESG concerns, market participants also face the risk of stranded assets. If tenants are demanding green buildings, there is a serious risk that the old types of buildings can no longer be rented out at "market" rents, and can only be rented out at deep discounts, the so-called brown discount. Ever stricter government regulation contributes to this risk, called *transition risk*. The next section digs deeper into this.

28.1.2 REGULATORY ISSUES: GOVERNMENT POLICIES FOR A GREEN REAL ESTATE SECTOR

Energy efficiency in real estate is and will remain a key policy issue for governments. According to most economists, the best action governments can take to increase energy efficiency in the real estate sector and beyond is to raise energy prices through a *carbon tax* or a *cap-and-trade system*, the most efficient and cost-effective ways of reducing greenhouse gas emissions. Essentially, the government makes CO₂ emissions expensive and then leaves it to market forces to decide on measures to cut emissions. The industry will then introduce the most cost-effective measures first, which are also most efficient for society as a whole. These measures can also change consumer behavior. For example, Europeans have traditionally driven much smaller cars than Americans because petrol was

⁴ Deloitte's new corporate headquarters in Amsterdam has been certified as the most sustainable office building in the world by BREEAM. The Edge was awarded an "Outstanding" rating with the highest-ever BREEAM score of 98.36 percent when it came on the market. It took the title from One Embankment Place in London.

⁵ Environmental, social, and governance (ESG) is shorthand for an investing principle that prioritizes environmental issues, social issues, and corporate governance. Investing with ESG considerations is sometimes referred to as responsible investing or, in more proactive cases, impact investing.

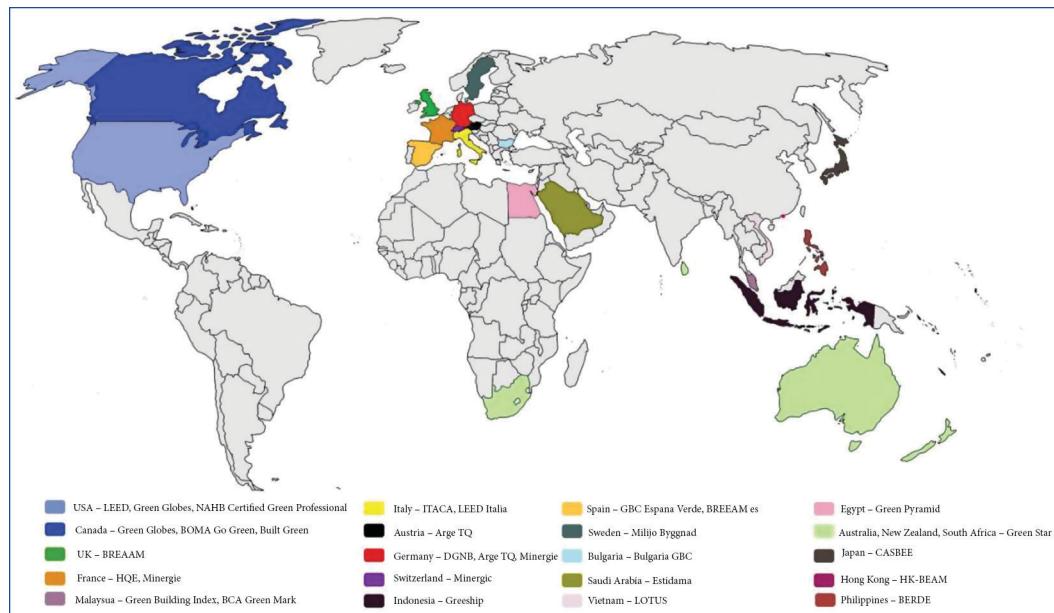


EXHIBIT 28-1 Green Building Labelling all Over the World.

Source: Mattoni et al., 2018.

much more expensive than in the U.S., incentivizing the purchase of more fuel-efficient cars. The same principle can be applied to energy consumption in buildings.⁶

A recent U.S. poll showed that the majority of American voters favor this type of approach, but most of their politicians do not. So, there is no U.S.-wide cap-and-trade system and no broad energy tax, although some states have implemented these policies. The EU has a cap-and-trade system, and an active emission trading system based on that. China introduced a system in 2021. But these systems are not as effective as they could be, since the emission caps tend to be set at levels that do not really bite, so large emitters are not strongly incentivized to change their ways. A truly global system with meaningful emission caps would clearly be more effective.

Governments are more likely to choose the carrot over the stick when it comes to CO₂ reduction policies: *subsidies* to stimulate green real estate investments abound. This may sound like a good idea, but it is doubtful whether these subsidies lead to outcomes that are optimal for society, and whether they are an effective way to spend tax money. Politically motivated schemes are often short-lived, changeable and/or unclear, whereas investments in green technology that have a long time to come to fruition would benefit from long-term clarity. For example, a heat pump subsidy fund opening on January 1 may be quickly used up for the year, and hamper consumer adoption, as many decide to wait a year for the subsidy to return. It also involves much higher bureaucratic and administrative costs. Subsidies can also be a waste of money, for example subsidizing actions the market would take on its own because of sufficient payoff without government stimulus.⁷

The third course taken by governments is to prescribe energy efficiency through *building codes*. As far back as the 1970s, spurred by the first energy crisis, these codes began to incorporate energy requirements such as insulation and double-glazing. Developers will usually go as far as the building code requires and no further, but the codes have steadily become stricter. New buildings are now

⁶ From a social justice view, a carbon tax could seem an inequitable, regressive tax, as the poor spend a larger portion of their income on energy costs and thus unfairly bear its brunt. However, funds raised by the measure could be redistributed to the poor to ensure the system operates equitably.

⁷ For example, switching from incandescent flood lights to LED floodlights has a payback of less than two years in terms of money saved on energy costs without the need for subsidies.

almost energy-neutral in many European countries and the U.S., though not yet in countries like China and India. The trend will spread further in the future.

However, building codes only apply to *new* construction. The problem is that we only build a few percent of the existing stock of buildings in any 1 year, so to upgrade the overall building stock we must also retrofit existing buildings. *Stricter energy performance criteria* are therefore also emerging for *existing* buildings. Examples are the minimum energy efficiency standards in the United Kingdom, France, and the Netherlands.

These existing and potential measures create transition risk or stranded asset risk for property investors. For example, as energy prices go up, tenants will prefer well-insulated buildings. And legislation is getting stricter all the time. If the real estate in question does not meet the regulatory requirements or measures up to what the consumer wants, its rental income is in danger and it will lose value.

Governments also stimulate energy efficiency through *energy performance transparency*, by introducing mandatory energy labels for buildings. Examples are Energy Star in the U.S., Europe's EPC label, and China's Green Building Evaluation Label. They tend to focus solely on the energy performance of buildings themselves, and not on other aspects of environmental performance. These labels subsequently also become the core of building codes and mandatory energy performance standards and give owners a yardstick on the transition risk of their buildings.

Another interesting tool in that regard is CRREM, designed specifically to map transition risk relative to the energy efficiency targets of the Paris climate agreement. Exhibit 28-2 provides a picture illustrating how the CRREM tool works. CRREM provides a pathway for a building or a property portfolio to the 2050 energy efficiency target resulting from the Paris goals, given expected technological development. The curved line in the graph is a stylized version of such a pathway. Starting at the building's current level of greenhouse gas intensity, an investor can determine how long it will take the building to perform worse than the pathway (to the right of the curve), after which transition risk starts getting salient. In this example, that would be in 2030, so a refurbishment before that year would be needed to avert that risk. Given the energy efficiency improvement resulting from that refurbishment, the tool then provides information regarding the time period for which the building is out of the risk zone, that is, to the left of the curve, helping investors to make the trade-off between the size of the refurbishment and their willingness to run transition risk. This tool is used by some of the leading institutional real estate owners globally.

However, the question is whether government involvement is necessary at all. Refurbishing an existing building to make it energy efficient is an investment, and if that creates financial value, the market should make the investment without government stimulus. We now turn to the question of whether energy efficiency investments do indeed add financial value.

28.2 FINANCIAL CONSIDERATIONS FOR OWNERS AND DEVELOPERS

28.2.1 DOES ENERGY EFFICIENCY ADD FINANCIAL VALUE?

Building owners and developers may decide to choose high levels of energy efficiency based on moral grounds, because they are worried about climate change. The question we aim to answer here is what they should do if they simply consider the effects on their own financial position.

In the case of new buildings, in many countries in Europe and in the U.S. this is not much of a choice since building codes make top energy performance mandatory. Even where this does not yet apply, it may make financial sense to target top energy performance. As we said before, the main goal of regulatory initiatives is to keep global warming below 2 degrees Celsius and try to limit it further to 1.5 degrees. However, with global warming now already at over 1 degree, it seems likely that regulation will get stricter in the coming decades. The implication for building owners is that they should not stop at *current* performance standards but should anticipate *future* standards in order to position their real estate assets favorably and avoid future transition risk. For example, an owner's building in 2025 will likely still be in existence in 2050, even if the owner has sold it long before then. If she sells in 2035, the building's value will partly depend on the prospective cash

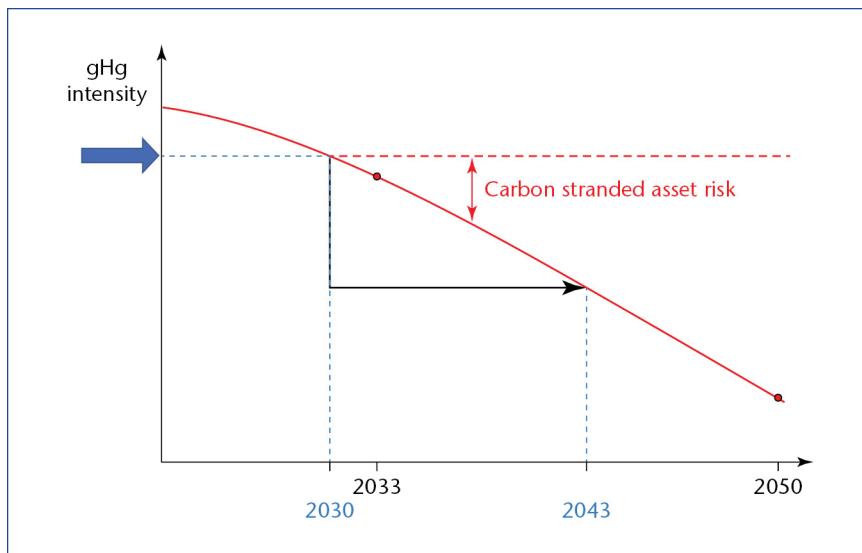


EXHIBIT 28-2 The CRREM Transition Risk Assessment Tool.

flow between 2035 and 2050 and by its prospective market value in 2050. With regulations getting stricter, the value of non-compliant buildings will suffer.

The next question is what owners should do with existing buildings. To decide whether to invest in energy efficiency improvements, the investor needs to look at the consequences of improved energy efficiency on the different aspects of a building's financial performance, in combination. For some reason, there seems to exist a misplaced market tendency to use the payback period for this. Building owners often argue that if the payback period of a green refurbishment is longer than their holding period, they do not invest in energy efficiency improvements.

This is of course the wrong way to evaluate investments, principally because it neglects all cash flow coming after the payback period, but also because it neglects the time value of money. The right tool for looking at the economic or financial value of energy efficient buildings and determining the added value of improvements is of course the net present value rule.

An investment in insulation measures or in a heat pump system creates different effects on cash flows and the discount rate. For example, if energy costs for tenants decrease, they may be willing to pay a higher rent. Or the government may decide to tax sustainable buildings more favorably. If it is easier to find a tenant for an energy efficient building, vacancy risk goes down, and so may the discount rate. If the refurbishing leads to higher building value, this value premium is also a key aspect of the investment framework. The net present value formula allows one to evaluate these different aspects of the investment, both individually and in combination.

Outcomes of the net present value rule are only as good as the information one puts into the formula. We therefore look at the available academic evidence regarding the effects of green investments: its costs; its effects on rents and occupancy rates; on energy costs and possible taxes; on risk, risk premiums and the cost of capital; and finally, on sales value.

28.2.2 THE EVIDENCE: CASH FLOWS, DISCOUNT RATES, VALUES

Starting around 2008, a broad international academic literature has emerged regarding the financial effects of energy efficiency investments in buildings.⁸ This has resulted in a considerable body

⁸ Miller et al. (2008), Eichholtz et al. (2010) and Fuerst and McAllister (2011) were the first to look at green value effects for offices, and Brounen and Kok (2011) for residential. The literature that has emerged since then has been reviewed in three useful meta-studies with good coverage: two for residential (Cespedes-Lopes et al., 2019; Zhang et al., 2018), and one for both commercial and residential (Dalton and Fuerst, 2018).

of academic evidence, but the strength of it varies. Regarding rents and sales values, the evidence is strong, but it is less convincing for costs and the discount rate. In addition, the evidence is not equally strong for the different property types: for offices and residential real estate, it is very broad and convincing, but for retail and logistics, let alone for ‘alternative’ commercial real estate sectors such as hotels, there is hardly any evidence to speak of. Third, the strength of the evidence varies a lot across countries. It is good in the U.S., some European countries, and some countries in Asia, but weak, or non-existent, in other regions. But this is only a matter of time: academics all over the world are investigating these issues, with evidence growing in quality and quantity.

First, we look at the **costs for green refurbishment**. Unfortunately, the literature on that topic is limited, with most papers providing case studies. Kok, Miller, and Morris (2013) do cover a larger sample of office retrofits between 2005 and 2010 and show that the capitalized benefits of light retrofits outweigh the costs. Another study providing solid statistical evidence in the U.K. (Chegut et al., 2019) shows an average marginal cost of 6.5 percent of green versus conventional refurbishment, going up to almost 30 percent for buildings with the highest green ratings with a similar payoff.⁹

For the effect of green refurbishments on **energy costs**, the challenge is in lease structures that pass through operating expenses to tenants such that any investments made by the landlord accrue to the tenant. To incentivize landlords, new green leases are needed that result in a rental payoff in exchange for energy investments paid for by the landlord. However, reduced energy costs do seem to affect tenants’ willingness to pay rent. Eichholtz, Kok and Quigley (2013) show an average rent increase for energy efficient buildings of 95 percent of the reduction in the energy bill, irrespective of the rent contract type. Tenants do not seem to care whether they are paying for energy or rent; they are looking at their total occupancy costs and are willing to pay more rent when the energy bill is lower and the quality of the space is higher.¹⁰

The evidence regarding **green real estate rents** is now fairly solid, with over 40 peer-reviewed papers published to date, covering different regional markets and phases of the cycle, and including both commercial and residential property. The evidence overwhelmingly points to rental premiums for green buildings. And the stronger the energy efficiency, the higher the rental premium. The first paper to document these rental premiums was Miller, Florance, and Spivey (2008), finding a 5 percent to 10 percent premium. These levels have been confirmed in subsequent literature, overall averaging 6 percent, with 8.2 percent for residential and 5.4 percent for commercial real estate (Dalton and Fuerst, 2018). On average, premiums for the U.S. do not differ significantly from those in the rest of the world.

The **discount rate** that needs to be employed in the net present value formula should reflect the risk of the investment. So, the key question here is whether green building design or features affect real estate risk. If green real estate is less risky, this reduces the required cost of capital and therefore also the discount rate. Here, the literature is small, but it tells a consistent story. First, green properties are more liquid than conventional ones, especially when selling in bad markets, where their time on the market is shorter (Brounen and Kok, 2011). They also have a higher and more stable occupancy rate (Eichholtz et al., 2010 and Kok et al., 2012). Furthermore, rents are more stable than energy prices, so if a building consumes less energy, there is less exposure to energy price volatility, implying less risk. At the portfolio level, all this results in a lower beta, the yardstick for systematic risk (Eichholtz et al., 2012). Furthermore, An and Pivo (2020) show that commercial mortgages on green buildings have a 34 percent lower chance of default.

⁹ When incorporating these percentages into the net present value equation, it is important to realize that they relate to the value of the structure only, and not to the land on which it is built. In contrast, the market value premiums associated with energy efficiency relate to the total building value, including land.

¹⁰ For example, more energy efficient buildings are also correlated with having more natural light, a feature cherished by most occupants. This has been borne out by a few studies. For example, Miller et al. (2009) show that over 50 percent of tenants felt they were more productive in green buildings. Unfortunately, professional productivity from better and more energy efficient environments remains difficult to measure.

This lower risk would suggest a lower cost of capital for green buildings, and that is indeed what the—very limited—literature suggests. For greener REITs, the interest rate costs are lower: the commercial mortgages on their green certified buildings have 24–29 basis points lower interest rates than those in their conventional buildings, and green REITs' bonds have lower spreads in the secondary markets (Eichholtz et al., 2019). An and Pivo (2020) also find lower spreads for mortgages on green commercial property, but they find a much smaller difference. Barron et al. (2019) show that the cost of equity for green REITs is 35–40 basis points lower than for conventional REITs.

Finally, we summarize the **sales value of the building**. Here also, the literature is quite strong, with evidence for different markets, time periods, and building types—mostly residential and office. Again, this literature points to a broad consensus: energy-efficient building command a sales premium, and that premium goes up with the energy performance of the building. Dalton and Fuerst (2018) survey this work in their meta-study and again show a clear consensus. Overall, they find an average sales premium of green certified versus conventional buildings of 7.6 percent, with a premium of 11.5 percent for commercial real estate and 5.5 percent for residential real estate. These results are based on past data. The effects are likely to become greater if energy consumption regulation gets stricter and /or energy gets more expensive, increasingly widening the difference between stranded and non-stranded assets.

We can employ this evidence to investigate the marginal value of a green building versus a conventional one using the net present value formula. When making reasonable assumptions regarding holding period, inflation, conventional building rents, and the risk-free rate, we tend to find positive marginal net present value effects. This suggests that energy efficiency investments in buildings should be made on the profit motive alone. However, the outcomes of such modelling exercises also suggest that going to the highest energy efficiency standards (i.e., net zero consumption of energy) is not adding sufficient financial value yet to be worth the cost, unless non-financial ESG objectives are brought into the equation.

Even with sufficient incentives to build greener new stock, we must keep in mind that the existing global building stock is vast, and turning around so much real estate would require huge capital and labor outlays. There is room for some optimism here: humankind has shown that the profit motive is a great mover, and that the opportunity to make money tends to remove obstacles that looked insurmountable before.

28.2.3 THE TIMING OF ENERGY EFFICIENCY INVESTMENTS

Investment decisions not only involve whether, but also when to deploy capital. Building owners looking at the information above may decide to act now, but it may be better financially to wait for improved and possibly cheaper green technology. Recall the discussion on option value from Chapter 16. The same rules apply here.

Exhibit 28-3 depicts a stylized picture of the benefits of energy efficiency investments, given a specific state of green building technology and cost, depicted by the top line in the graph. This suggests that capital employment at first yields a relatively large reduction in energy use, and subsequently a declining marginal benefit per dollar invested. An investor can reap a large part of the total possible benefits in energy efficiency by deploying a limited part of her capital. Somewhere on that curve lies the point where the investment improves a building's energy performance—and therefore the rent, discount rate and value—just enough for the investment to yield a net present value of zero. This point is associated with an investment level of I_1 and an energy performance level of E_1 . All investments lying to the left of that point are NPV-positive and should therefore be undertaken.

If green building technology improves and gets cheaper per kWh saved, the curve will shift downwards, increasing the bang for every buck of investment, as the resulting energy use drops from E_1 to E_{1a} . That is the direct effect of technological improvement on energy efficiency. But the graph shows that there is an indirect effect on top of that. The decrease in energy use will have beneficial effects on rental income, the discount rate, and building value, which implies that the level of

investment that was NPV-zero before the change in technology will now be NPV-positive. So, the NPV-zero point will shift to the right along the new curve, making a greater range of investments financially viable, up to investment level I_2 . That creates yet more impact in terms of energy usage reduction, with energy use dropping further to E_2 . Technological progress has a double effect on the impact of energy efficiency investment: it increases the efficiency of a given investment, and it invites more value-adding investment.

With continuing technological progress, it will get ever cheaper to reach a given level of energy efficiency, and the NPV-positive range will get larger and larger as the NPV-zero point continues to shift to the right. An increase in energy prices, whether caused by market forces or by government actions, would have a similar effect.

Investors may thus argue that to deploy their capital more effectively, they should wait for technological progress to happen, leading to a higher net present value in the future: future present value. However, waiting means foregoing immediately improved cash flows from green refurbishment, thanks to higher rents and lower vacancy rates. The bigger these advantages, the larger the resulting opportunity loss over these years. Also, the future present value needs to be discounted to calculate the present net present value. This will occur at the higher, non-green discount rate. So what to do?

Fortunately, net present value thinking also provides clear guidance in this instance. Waiting creates an opportunity loss, the size of which is measured by the opportunity cost of capital, which is the discount rate used in the net present value model. The gain of waiting is measured by the speed of technological change, here approximated by the speed of change in energy reduction accomplished by a given level of investment. If the discount rate is higher than the speed of technological change, then it is better to start investing as soon as possible.

For example, with a discount rate of 5 percent, a 3-year wait would require at least a 16 percent higher future present value to obtain a positive net present value and make the investor better off now. In other words, the technology needs to get 16 percent cheaper – or more effective – in these three years, and if it does not, holding off on the investment is not a good idea. If LED flood light

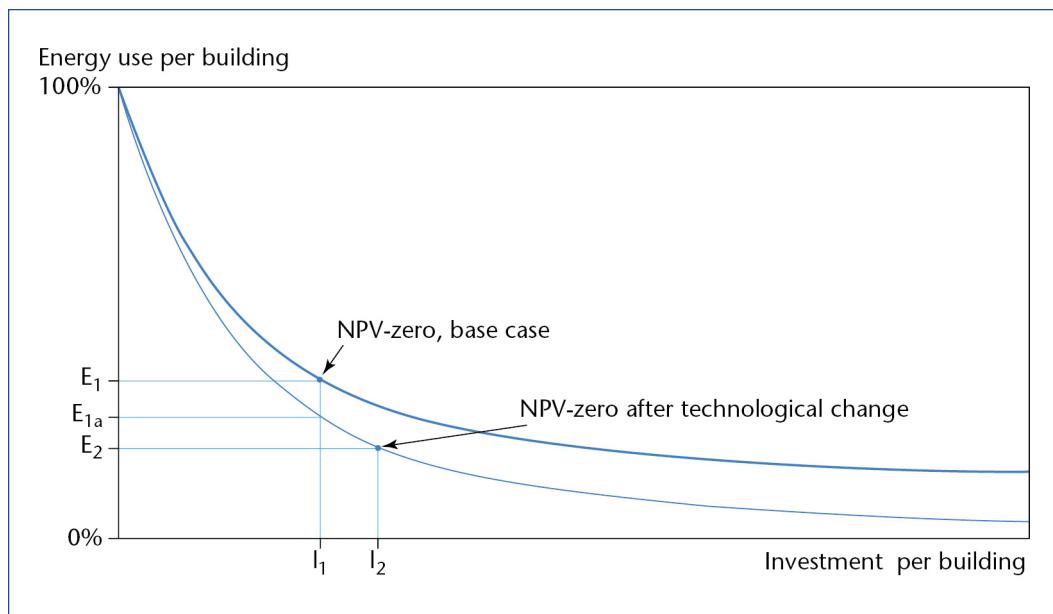


EXHIBIT 28-3 Investment Impact in Energy Efficiency with Technological Change.

costs \$10 now but will likely cost only \$3 within two years, waiting on replacement would make sense since the speed of the cost reduction is so much higher than the discount rate.

To use this rule in practice, it is important to realize that many green refurbishment measures are already mature and that all of them, from insulation to solar panels, involve substantial labor costs. For example, in solar panels, typically one-third of today's costs may be installation labor with low productivity growth, and these costs are going up with wage inflation. The next one-third involves low-tech auxiliaries such as electrical cables, inverters, and the hardware to connect the panels to the building, of which the costs also tend to go up with inflation. Only the final third involves the actual panels. And of that, not all money goes to the producer in China, but also to transportation and the profit margin of the wholesaler in the country where they are installed. Overall, the production cost of the panel itself may be just 10–20 percent of the total investment cost. So even if prices at the solar panel factory would come down by a lot and the efficiency of the panels would go up, the other components of the bill would lessen the speed of change in energy reduction accomplished by a given level of investment. Similar reasoning holds for other components of green refurbishment.

As a final word on the timing of green refurbishments, it is important to remember that every building has a natural maintenance/redevelopment cycle. For a large refurbishment, just getting into the building while accommodating the sitting tenants is a major part of the overall investment cost. A building owner should work within the cycle to keep costs low. However, the cycle can also be adjusted. Building owners who are mindful of rental contract expiry dates can employ temporary vacancy as refurbishment opportunities, possibly even merit bringing a renovation forward a few years.

28.3 PHYSICAL CLIMATE RISK AND REAL ESTATE

28.3.1 LOCATION, LOCATION, LOCATION

“Location,” location, location, it holds for real estate as much as it does for physical climate risk, and if your real estate happens to be located in the spot where disaster hits, there is no escape. This makes physical climate risk particularly salient for real estate, yet it has been all but ignored by the industry until quite recently. By and large, it was not part of the risk assessment of commercial real estate financiers, it did not play a role in real estate valuation models, and investors did not include it as a criterion in their investment decisions. However, this is changing rapidly, with the incidence of physical climate risk mounting all over the world.

The nature of climate risk is both *specific* and *systematic*—partly diversifiable and partly undiversifiable. It is specific since it can play out at the *micro* location level. For example, much real estate has been built near water along coasts and at river mouths, following the commercial logic of the past. Climate change brings typhoons and hurricanes, rivers and seas flood and buildings are vulnerable, stuck in their locations. But the effects can be hugely different across buildings, even over very small distances. One building may be flooded and another one just 100 metres away unaffected. If wildfires break out, houses may burn down on one side of the valley but not on the other. Geographical diversification is key to managing this risk at the portfolio level.

However, that brings an investor only so far. Once all the diversification benefits are reaped, physical climate risk still adds another layer of risk to real estate investment, which is the systematic component of that risk. Global property damage is going up, partly due to climate change. Much of that damage is still insurable, but the increased insurance bill will land on the desks of the property owners, and they will also bear the risks that are not insurable.

Most forms of physical climate risks are already noticeable now, with the temperature only just over 1 degree Celsius higher than Paris accord guidelines, and further increases are certain. It is also certain that sea levels will rise, even if we can keep the temperature within the limits of the Paris

agreement; hurricanes will get more frequent and more severe and visit higher-latitude locations; rivers will flood more often, and there will be droughts, wildfires, heat stress, and soil subsidence.

Both the frequency and the impact of these climate events will be higher in emerging economies. Physical effects of rising temperatures will be worse around the equator, where greater poverty and more non-air-conditioned open housing will make heat issues harder for the population to bear. Furthermore, the infrastructure is less resilient, making the impact of a given climate event greater. There is less physical protection in less developed nations, for example weak dikes against floods in Bangladesh, compared to strong dikes in the Netherlands. It is also harder to reach disaster locations to provide relief, with poor roads and other transport infrastructure.

Good governance can reduce the impact of climate events, but this is weaker in emerging economies. Urban planning could prevent building in flood zones or locations with likely soil subsidence, but in emerging economies, people are more likely to build in unplanned ways. There is also less government money available for relief when disaster strikes and for the repair of critical infrastructure such as roads and bridges.

Since the mid-2010s, good-quality academic research has emerged on the degree to which these risks are getting priced in real estate markets. The availability of the evidence differs across risk types and property types, as this is still an emerging field of research.¹¹ The evidence is most abundant for housing, and less so for commercial real estate; more for (coastal) flood risk than for the other risk types; and more for the U.S. than for other countries. The consensus emerging from that literature is that climate risk discounts exist, and that they are rising. Higher property insurance costs are being capitalized into lower property value. However, a debate is going on concerning the question of whether the property discounts are already large enough. Many papers find that the positive amenity value of living close to the sea or a river, is still higher than the climate risk discount.¹²

28.3.2 PHYSICAL CLIMATE RISK: THE BIG SIX

The major effects of climate change are a rise in sea levels, river flooding, more and worse hurricanes, wildfires, heat stress and finally, soil subsidence. These are happening already.

Rising sea levels are inevitable even if the world manages to stay within the Paris goals. Seawater expands as it warms and melting land-based ice caps and glaciers further increase the total body of water. Although all these together—including Antarctica, Greenland, and global mountainous regions such as Tibet and South America—make up just 2.1 percent of all the earth’s water, if all of this ice were to melt, sea levels would rise by 230 feet, according to the U.S. Geological Survey. This would wipe out most of the state of Florida and much of the inhabited coastal land of many nations like Bangladesh, Nauru, Indonesia, China, Maldives, Vietnam, the Netherlands, and others. However, that scenario is unlikely, and it will unfold very slowly. But it is not impossible, if global temperatures keep on rising.

What is going to happen with high probability and impact is a limited increase in global sea levels. The sea level has already risen by an average of 35 centimeters globally. Even if CO₂ emissions are brought down in accordance with the Paris agreement, the sea-level rise will still theoretically be 40-100 cm by 2100 and at least 3 feet higher in 2300. In the longer term, sea levels will go on rising, and will remain high for millennia. This will increase the frequency of high-tide coastal flooding.

¹¹ Contat et al. (2024) wrote the first review of the literature in this emerging field. They look at pricing effects, loan performance and migration effects, focusing on sea-level rise, flooding and wildfire risk.

¹² See, for example, Gourevitch et al. (2023), Niu et al. (2023), and Gibson and Mullins (2020) for reflections on amenity value versus the size of the climate risk discount, and Addoum et al. (2023) and Rehse et al. (2019) about pricing and liquidity effects in U.S. commercial property. Miller, Gabe and Sklarz (2019) found that coastal discounts dissipate after one or two years.

The Intergovernmental Panel on Climate Change (IPCC) calculates that at least 10 percent of the global population is living in low-elevation coastal zones exposed to the risk of high-tide flooding. The question is whether affected cities will still be livable. Venetians, for example, have experienced frequent occasional flooding, and they hardly live in the flooded parts of the city anymore.

The incidence of **river flooding** is increasing, with changed rain patterns and behavior of rivers due to climate change. On the one hand, the summers get drier in some areas, but the rain that does fall tends to get more concentrated and severe: low-pressure systems form above warmer seas and oceans, providing more perspiration, and warmer air can hold more moisture, so rains are getting heavier. Many riverbeds are currently not suited for this. Rivers are increasingly fed by rain only, rather than by gradually melting snow and mountain glaciers. This makes water levels more unpredictable, possibly hampering river transportation in dry periods. Yet at the same time, summer flooding of rivers is occurring more often and with greater severity throughout the world. All this requires better water management, to enable conservation of river water and its release in dry periods.

Historically, cities often grew at the mouths of big rivers for reasons of commerce, and many of the world's major harbors today are at river mouths (think Shanghai, Rotterdam, New York). But these rivers need to release their water to the sea, and rising sea level will make that increasingly difficult. Sea level rise is particularly salient through increased river flood risk in such cities.

Hurricanes (or typhoons) can only form and grow if the ocean surface temperature is higher than 26.5 degrees Celsius (79.7 Fahrenheit). The higher the temperature, the more water evaporates, creating the 'engine' for the hurricane to build up. Thus, rising sea temperatures lead to more frequent, more severe, and longer-lasting hurricanes. They also create a longer hurricane season, and as the sea temperature rises at higher latitudes, new locations are affected. Hurricane Sandy (2012) was an example of that. Helene (2024), with a 500-mile pathway of destruction, was an example of a storm with a much broader path of damage than most previous hurricanes, causing over \$50 U.S. billion in damage and killing more than 200 people.

Even with CO₂ levels meeting the Paris targets, hurricane risks will increase, both in terms of probability and severity. As the sea surface temperatures go up, hurricanes near the U.S. become more frequent and longer, migrate north, and become more severe. The same holds for typhoons in Asia. Increased severity means more damage, with the average now at around \$10 billion per hurricane. These effects are inevitable and will get more severe. There will be other Katrinas, Sandys, Ians and Helenes, with all the damage they inflict. Not just in the U.S., but also in other parts of the world.

In dry areas, **wildfires** are happening with unprecedented frequency and in new locations, due to longer, more severe drought periods. Undergrowth dies out and trees die and then dry out completely. This creates a tinderbox effect. Forests are no longer green and wet but dry and flammable, meaning more and worse fires. Where a couple of farms might have burned down in the past, this could now be whole neighborhoods, and wildfires may even threaten larger residential areas. The 2025 Los Angeles fires were a grim example of this, and this is also seen in other parts of the world. Sydney, Australia, for instance, is now under threat of repeat fires, as are parts of South Asia, with bush fires seen in Malaysia.

Wildfires are also beginning to show up in new places, even including chilly Sweden and Russia and the traditionally wet Netherlands and Belgium, among other places previously considered low-risk. Again, these events are already happening and will get more severe.

Another risk affecting city dwellers and their buildings is **heat stress**, as summers are becoming hotter and longer. Concrete and brick have a very high capacity to store heat by day and radiate it at night, which creates urban heat islands. Humans can cope with daytime heat, but they need to cool off at night. If they are unable to do so, especially in a city with no air conditioning, and if this is not just for a few days but for weeks on end, it affects people's health. The vulnerable and less healthy elderly

and poor are hardest hit. For example, the 2003 heatwave hitting France resulted in excess mortality in Paris that was an order of magnitude larger than it was in the rest of the country. Because of this, cities may become unlivable in certain parts of the world, mostly in emerging economies. In addition, cities that did not need air conditioning before will now adopt it en masse, leading to higher energy use.

The emergence of unlivable cities will affect urbanization. In Europe and the U.S., the urbanization rate is above 75 percent whereas in India, for example, it is just under 40 percent. Under normal circumstances, urbanization would continue, creating more productivity and prosperity. But if cities become unlivable, because of the heat island effect, urbanization may slow down, and economic growth with it.

Finally, **soil subsidence risks** occur as drought creates unstable soil and destabilizes buildings, for example where cities are built on mud, peat or clay. This is often the case near the coastline and would apply to most cities built at the mouths of river deltas. Building damage gradually increases, and cities can begin to sink, which is especially problematic where they already face sea flooding.¹³ While little academic evidence is yet available worldwide, the importance of subsidence has already been flagged in Indonesia, where Jakarta is sinking. The plan is therefore to move the seat of the national government from Jakarta on Java and relocate it to Sumatra. In the Netherlands and other river deltas, it may also become necessary to invest many billions to make buildings in cities more resilient.

28.3.3 EFFECTS, RISKS, AND INSURANCE

Having outlined these six main physical risk categories, the question is how these risks can be managed. For risks that possibly lead to physical property damage, the standard approach is to take out property insurance. However, not all climate risks are insurable, and for those that can currently be insured, it is not certain whether this will remain the case. To shed more light on this, it is useful to distinguish between first, second and third order effects of physical climate risk, ranging from direct to indirect effects; from immediate to longer term, and from relatively small to very big.

First order effects are immediately apparent and include both physical damage and temporary hindrance to life and business. They might be expensive, but they will not affect the value of buildings very much or at all. Restoration is still doable, and damages are currently still largely insurable, even while insurance premiums may have to go up to cover the increased risks and damages. Property owners may also want to invest in adaptation measures, for example, by relocating their building systems hardware from the basement to a higher floor to improve flood resilience, or by increasing insulation to prevent heat stress. The net present value framework can be used to analyze the financial added value of such resilience investments that (partly) avert the above trends.

Second order effects are medium-term and may impact property valuations. This could happen if, for example, floods repeatedly damage the buildings in the same areas, or wildfires keep burning down the same towns. This makes these buildings and locations less attractive, possibly leading to a decreasing local demand for space, and increased vacancy. This will affect cash flows and will lead to losses in value, and it may have negative regional economic spillover effects. The degree to which this will happen depends on the balance between the amenity value of a location and the increased climate risk. For property owners, the ways to manage this are regional diversification, and to sell buildings in exposed areas.

To make thorough hold-sell analyses, building owners should employ the net present value framework, incorporating the lower rents or lower occupancy assumptions over time at high-risk locations. Vacancy will be higher due to tenants moving out and to damage repair periods. There

¹³ Ohenhen et al. (2024) provide an analysis of the interplay between soil subsidence and sea-level rise for U.S. coastal cities and calculate the number of people and the aggregated value of the properties at risk. They also discuss some adaptation measures.

will be higher insurance premiums and more frequent and higher capital expenditures for repairs. The investors' discount rates will move up, due to higher equity risk premiums and mortgage rates, and terminal values will be lower as exposed buildings are less attractive on the market. Again, the magnitudes of many of these changes are already documented in academic research, helping developers and building owners make informed decisions. Besides that, a climate risk information industry is emerging.

TEXT BOX 28.2 CLIMATE RISK ASSESSMENT TOOLS

How do we measure climate risk and relate it to specific locations and buildings? It starts with climate risk maps, which are available for more and more locations, though usually for developed countries and less so for emerging markets. In the U.S., FEMA (the Federal Emergency Management Agency) produces maps for the different physical climate risk types such as flood risk and wildfire risk. Governments of other countries provide similar services.

However, it is not trivial to translate these maps into a clear risk score for a building or a location, for a number of reasons. First, the disaggregation level of these maps is not always very deep. Second, the big-six physical risk types may be correlated. For example, land subsidence in river deltas may happen at the same time as sea level rise, creating a larger combined risk. Third, these maps tend to be created for engineering purposes and are not based on actual probabilities of physical risk events, but tend to err on the safe side, reflecting a kind of safety buffer that is best practice for civil engineers, but less useful for building owners who want to have an accurate picture of the risk they run. Fourth, the maps look at the climate risks given the state of nature today, rather than using predictions about the future state of nature. Last, they are not always updated very regularly.

To support developers and building owners, private companies have therefore created climate risk scores at the building or micro-location level. Some of these companies are established financial information providers, such as Moody's, SwissRE, S&P Global, MSCI, and Bloomberg. The field also contains companies that focus on climate risk alone, such as ClimateCheck, FirstStreet Foundation, Jupiter Intelligence, and Sustainalytics.

These initiatives clearly provide a step in the right direction, but climate risk information remains an emerging industry. Most problematic is that different climate risk services provide wildly different scores for a given location. The cross-sectional correlation between the risk scores of the different providers is close to zero, casting doubt on the quality of the information. Clearly, there is a lot of room for improvement.

The increased incidence of climate risk may also affect the extent to which property insurance remains available. For risks to be insured, they should typically be low probability—high impact, preferably with a well-understood probability and a predictable impact. Due to climate change, risk probabilities are changing—making them less predictable. Incidence may go from low probability to high probability in certain locations. At a global or national level, this does not seem to be problematic yet. While annual damages go up, this is only partly due to climate change, but more to inflation and (urban) economic growth: there is simply more economic value at risk. However, at a local level, insurance is starting to become an issue. For example, several insurance companies had pulled out of California before the 2025 fires, citing the high risk of wildfires as a principal reason to do so.

Physical climate risk has a twofold effect on insurance premiums. First, the increased frequency of damaging climate events, as well as the enhanced severity of a given event, may directly lead to soaring insurance premiums. Second, when insurance companies withdraw from certain locations, the local insurance market will get less competitive, also leading to higher premiums. Ultimately, this

may make insurance unattractive, and if this shift is great enough, property damage becomes uninsurable. The larger commercial property owners are responding to this by turning to self-insurance, where they use their own balance sheet as a buffer to pay for property damages. Of course, this is only feasible for property companies with a portfolio that is regionally diversified.

For (regional) commercial property owners that are too small for this, the question is whether real estate can still be financed if commercial property insurance becomes prohibitively expensive or ceases to be available. A standard feature of any mortgage contract is collateral property insurance. This is so common that it can easily be overlooked, yet without it, property financing as we know it would likely come to a halt. Currently, maximum mortgage amounts take into account the value of both the land and the structure. If it would not be possible to insure the structure, then the maximum loan-to-value would just be based on the land (site) value, making the financing of (re-)development all but impossible, and with that, destroying the feasibility of the local property market.¹⁴ Of course, the government may step in and provide insurance from the public purse where insurance companies pull out. But that would keep people and businesses in locations where they'd better not stay. Moreover, even the public purse is limited. Would it be politically sustainable if the bulk of federal insurance money keeps going to owners of beachfront property in Florida?¹⁵

Third order effects are long-term and very large scale. If climate events keep hitting a region, this may result in severe deterioration of local economies and their real estate markets. Some regions may get hit by climate shocks so frequently that large-scale relocation of people and businesses will result. Businesses may decide they have "had enough" and relocate. Citizens may move away when they become tired of removing mud from their homes or apartments, replacing furniture every year. Such relocations may have a huge, long-run impact on real estate markets, making properties in affected location essentially worthless. It also implies that specific locations will fail, while other locations will rise, since the cumulative demand for real estate will likely not be affected. These are very long-run developments, possibly playing out well into the end of this century and beyond. Most of the world's population is not taking climate risk that seriously yet, but eventually they will. We will likely know more by the 5th edition of this book.

28.4 CHAPTER SUMMARY

Real estate can play a key role in reaching the Paris climate goals. For that, very high capital outlays are needed, especially in refurbishment of the existing stock. The net present value framework is the preferred tool to evaluate these investments, and ongoing academic research provides the required inputs.

Model outcomes based on these inputs generally show that energy efficiency measures tend to be value-enhancing. Greener real estate is, for the most part, a good business opportunity. Delaying only makes sense when the percentage increase in the per-dollar effectiveness of refurbishment is higher than the opportunity cost of capital. Even when approached purely from financial self-interest, sustainability investments add value and can be profitable, having a positive net present value. Owners who forego the investments in energy efficiency expose themselves to unnecessary transition risk from changing regulations and tenant preferences. With increasing regulation comes increasing urgency.

Even if climate risk mitigation investments accelerate, real estate will still be exposed to physical climate risk: rising sea levels and coastal flooding, river floods, hurricanes and typhoons, wildfires,

¹⁴ Holtermans et al. (2024) show that commercial mortgage delinquency rates associated with climate events are not just affected by declining collateral values, but even more so by the inability for borrowers to meet their payment obligations. This could be due to higher post-disaster capital expenditure, increased insurance premiums, and relocating tenants.

¹⁵ The U.S. National Flood Insurance Program is already running at a substantial loss with many billions at risk.

heat stress and soil subsidence. The incidence of these risk types is highly location-specific, and “immobile” real estate cannot escape. These risks cannot be completely diversified away. Physical climate risk adds a new layer of systematic risk to the real estate markets. Academic evidence shows that these physical climate risks are increasingly being priced in the real estate markets, but the strength of the evidence differs across risk types and markets.

Not all physical climate risk can be insured. While first order effects of physical damage can be insured for now, higher-order effects that largely impact regional demand for space cannot. On top of that, as the incidence of climate trouble increases, local property insurance markets may break down, which could have major consequences for the way real estate will be financed in these locations. It is unlikely that mortgage providers will regard an uninsured building as good collateral for a loan.

Finally, how far should a building owner go in striving towards greater energy efficiency? While better energy performance tends to be value-enhancing, making each building energy-neutral is likely not yet generating a positive net present value. On the other hand, the commercial real estate industry is among the sectors most exposed to physical climate risk. So, should building owners go further towards net zero?

The analogy of dealing with slow moving climate risk is the “tragedy of the commons,” where the village’s common ground is wrecked when self-interested individual farmers have their sheep over-graze and destroy the common ground, making everybody worse off in the end. To avoid this tragedy for the global real estate industry, pressing on to net zero and thereby helping the planet to stay livable for everyone seems like a good idea, with thriving future real estate markets as a result.

KEY TERMS

- Climate risk
- Transition risk
- Physical climate risk
- Flood risk
- Heat stress
- Urban heat island
- Sea level rise
- Soil subsidence
- Wildfire risk
- Intergovernmental Panel on Climate Change (IPCC)
- Paris climate goals
- Climate risk maps
- Environmental, social, governance (ESG)
- Energy efficiency
- Scope-3 emissions
- Green buildings
- Green refurbishment
- Energy performance labels
- Carbon tax
- Cap-and-trade system
- Building codes
- Energy performance standards

STUDY QUESTIONS

Conceptual Questions

- 28.1 Define transition risk and explain how it affects the real estate industry.
- 28.2 What are the three primary ways governments regulate or incentivize energy efficiency in real estate? Provide an example for each.
- 28.3 Why is the net present value (NPV) framework more suitable than the payback period for evaluating energy efficiency investments?
- 28.4 Explain how energy-efficient buildings can lead to financial benefits for owners and tenants.
- 28.5 What are the six main types of physical climate risks identified in the chapter? Provide an example of how each risk can impact real estate.
- 28.6 How do green certifications like LEED and GRESB contribute to improving energy efficiency in buildings?
- 28.7 Discuss the concept of “stranded assets” in the context of real estate and climate risk.
- 28.8 Why are some physical climate risks considered systematic and others specific? How does this distinction affect real estate investment strategies?
- 28.9 Explain the challenges associated with measuring climate risks at the property level. How are these challenges being addressed?
- 28.10 Under what conditions does delaying energy efficiency investments make financial sense?

REFERENCES AND ADDITIONAL READING

Addoum, J., P. Eichholtz, E. Steiner and E. Yönder. 2023. Climate Change and Commercial Real Estate: Evidence from Hurricane Sandy, *Real Estate Economics* 52(3), 687–713.

An, X. and G. Pivo. 2020. Green Buildings in Commercial Mortgage-Backed Securities: The Effect of LEED and Energy Star Certification on Default Risk and Loan Terms, *Real Estate Economics* 48(1), 7–42.

Barron, P., P. Eichholtz and E. Yönder. 2018. REIT Performance and the Cost of Equity, *The Routledge REIT Research Handbook*, 1st edition. Routledge.

Brounen, D. and N. Kok. 2011. On the Economics of Energy Labels in the Housing Market. *Journal of Environmental Economics and Management* 62(2), 166–179.

Cespedes-Lopez, M.F., R.T. Mora-Garcia, V.R. Perez-Sanchez and J.C. Perez-Sanchez. 2019. Meta-analysis of Price Premiums in Housing with Energy Performance Certificates (EPC), *Sustainability* 11(22).

Contat, J., C. Hopkins, L. Mejia and M. Suandi. 2024. When Climate Meets Real Estate: A Survey of the Literature, *Real Estate Economics* 52(3), 618–659.

Dalton, B. and F. Fuerst. 2018. The “Green Value” Proposition in Real Estate, *Routledge Handbook of Sustainable Real Estate*, Chapter 32. Routledge. London.

Eichholtz, P., R. Holtermans, N. Kok E. and Yönder. 2019. Environmental Performance and the Cost of Debt: Evidence from Commercial Mortgages and REIT Bonds, *Journal of Banking and Finance* 102, 19–32.

Eichholtz, P., N. Kok and J.M. Quigley. 2010. Doing Well by Doing Good? Green Office Buildings, *American Economic Review* 100 (5), 2492–2509.

Eichholtz, P., N. Kok and J.M. Quigley. 2013. The Economics of Green Building, *The Review of Economics and Statistics* 95(1), 50–63.

Eichholtz, P., N. Kok E. and Yönder.. 2012. Portfolio Greenness and the Performance of REITs, *Journal of International Money and Finance* 31(7), 1911–1929.

Fuerst, F. and P. McAllister. 2011. Green Noise or Green Value? Measuring the Effects of Environmental Certification on Office Value, *Real Estate Economics* 39(1), 45–69.

Gibson, M. and J. Mullins. 2020. Climate Risk and Beliefs in New York Floodplains, *Journal of the Association of Environmental and Resource Economists* 7(6).

Gourevitch, J., C. Kousky, Y. Liao, C. Nolte, A. Pollack, J. Porter and J. Weill. 2023. Unpriced Climate Risk and the Potential Consequences of Overvaluation in U.S. Housing Markets, *Nature Climate Change* 13(3).

Holtermans, R., Kahn, M. and Kok, N. 2024. Climate Risk and Commercial Mortgage Delinquency, *Journal of Regional Science* 64(4), 994–1037.

Kok, N., N. Miller and P. Morris. 2012. The Economics of Green Retrofits, *Journal of Sustainable Real Estate* 4(1).

Mattoni, B., C. Guattari, L. Evangelista, F. Bisegna, P. Gori and F. Asdrubali. 2018. Critical Review and Methodological Approach to Evaluate the Differences among International Green Building Rating Tools, *Renewable and Sustainable Energy Reviews* 82(8), 950–960.

Miller, N., J. Gabe and M. Sklarz. 2019. The Impact of Waterfront Location on Residential Home Values Considering Flood Risks, *Journal of Sustainable Real Estate* 11(1).

Miller, N., D. Pogue, Q. D. Gough and S.M. Davis. 2009. Green Buildings and Productivity, *Journal of Sustainable Real Estate* 1(1).

Miller, N. J. Spivey and A. Florance. 2008. Does Green Pay Off? *Journal of Real Estate Portfolio Management* 14(4).

Niu, D., P. Eichholtz and N. Kok. 2023. Asymmetric Information Provision and Flood Risk Salience, *MIT Center for Real Estate Research Paper* 23, no. 20.

Ohnenhen, L., M. Shirzaei, C. Ojha, S. Sherpa and R. Nicholls. 2024. Disappearing Cities on U.S. Coasts, *Nature* 627 (8002).

Rehse, D., R. Riordan, N. Rottke and J. Zietz. 2019. The Effects of Uncertainty on Market Liquidity: Evidence from Hurricane Sandy, *Journal of Financial Economics* 134(2), 318–332.

Zhang, L., J. Wu and H. Liu. 2018. Turning Green into Gold: A Review on the Economics of Green Buildings, *Journal of Cleaner Production* 172, 2234–2245.

29 Technology and Real Estate

A Synthesis

CHAPTER OUTLINE

- 29.1 Technology And Real Estate: A Synthesis
 - 29.1.1 Introduction to the Intersection of Technology and Real Estate
 - 29.1.2 Technology Applied to Design, Construction and the Metaverse
 - 29.1.3 Technology Applied to Data and Data Analysis
 - 29.1.4 Technology Applied to Marketing
 - 29.1.5 Technology Applied to Occupancy
 - 29.1.6 Technology Applied to Building Management Systems and Operations
 - 29.1.7 Technology Applied to Measurement and Analysis
 - 29.1.8 Technology Applied to Ownership Forms
 - 29.1.9 Technology Applied to Transaction Management and Reporting
 - 29.1.10 The power of multiple gradual improvements in Proptech and ease of integration
 - 29.1.11 The merits of mixed-skills teams in data-rich real estate investments
- 29.2 Chapter Summary

LEARNING OBJECTIVES

Readers of this chapter should understand:

- What is meant by proptech?
- How does proptech affect the real estate industry broadly and more specifically.
- Examples of proptech applied to data analysis, operations, asset management and investor reporting.
- Types of data available for real estate analysis.
- How better data has made the market more efficient.
- What do Building Management Systems do?
- What is Tokenization and Crowdfunding?
- How Tokenization has been facilitated by Blockchain technology.
- How appraisals are becoming more automated.

29.1 TECHNOLOGY AND REAL ESTATE: A SYNTHESIS

29.1.1 INTRODUCTION TO THE INTERSECTION OF TECHNOLOGY AND REAL ESTATE

The term “Proptech” might be seen as a tautology when we consider the essence of property itself. Property is fundamentally a form of technology—a primal innovation that shelters us from the elements, safeguards against predators, and provides a foundation for human survival. It is the original platform upon which layers of other technologies have been built to enable production, consumption, and culture.

What we often mean by “Proptech” today, however, is the infusion of information technology into this age-old concept. It signifies not property itself as technology, but the innovations transforming how we interact with, manage, and utilize property in an era defined by data, connectivity, and automation. Thus, Proptech represents a shift in perspective: not a new technology in essence, but a new technological layer over an ancient foundation.

Earlier chapters provide an introduction to real estate finance, focusing on how general finance theories and concepts can be applied to an asset class that is defined by its distinctions from mainstream

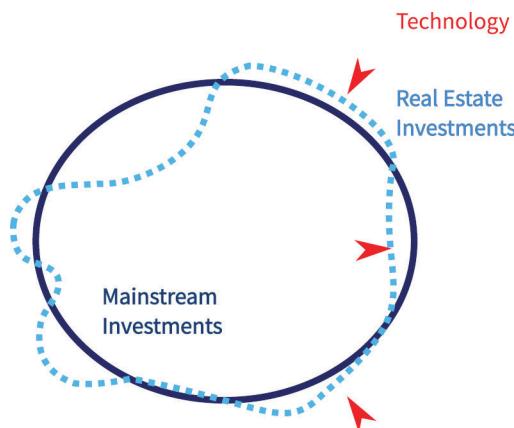


EXHIBIT 29-1 Proptech Reduces Frictions in Real Estate Markets, Gradually Aligning Real Estate with More Mainstream Investments.

investments like stocks and bonds. Real estate investments are uniquely shaped by a variety of frictions: high transaction costs, incomplete information, asset indivisibility, heterogeneity, the dual nature of residential property as both an investment and a consumption good, high operational costs, significant capital intensity, large environmental and social impacts, and intricate regulation and oversight.

Proptech emerges as a transformative force in addressing these frictions, gradually aligning real estate with the characteristics of more mainstream investments (Exhibit 29-1). However, the true potential of Proptech lies not just in making real estate more like stocks or bonds, but in enhancing its unique qualities—creating “better” real estate and fostering more complete, efficient, and equitable property markets.

Technological innovations in property (like the introduction of electric lighting, elevators or air conditioning) are nothing new. Developers, investors, operators and users of buildings are constantly striving to make the built environment more efficient, so that it requires less labor, energy, time, data collection—less of all inputs—while producing equal or better outcomes. The results of applying new technologies to real estate may be cost or time savings or revenue enhancing, thereby affecting values. Property designs may also be affected by technology as well as the physical operational facilities. As new technologies become more widespread in their adoption across the economy, it raises the bar for the property industry. Those engaged in the real estate industry must monitor new technologies as best we can and not fall behind.

The term “Proptech” is relatively new, gaining wider usage since maybe 2010. It likely draws inspiration from similar terms like “Fintech” (technology applied to finance), “Climatetech,” and other variations of “tech.” However, many companies involved in these technological applications to real estate are not new. For years, they have been collecting and organizing public and private property data and transaction records.

The Proptech umbrella spans a wide range of work functions, encompassing, but not limited to, the following categories:

- Design, Building Information Modeling (BIM) and Digital Twins
- Project Contracting and Management
- Property and Transaction Data and Analysis
- Financing Analysis, Sourcing and Support
- Equity Provision, Crowdfunding, Tokenization

- Financial Analysis and Valuation
- Brokerage and Leasing
- Development Support
- Marketing, Leasing and Sales
- Transaction and Title Management
- Occupancy management efficiency systems
- Building Management Systems and Operations
- Investor Reporting
- Energy Measurement, Management, and Wellness
- ESG (Environmental Social and Governance) Reporting
- Climate Risk Data and Analysis (Flood, Fire, Weather)

It is likely that there is an optimal point for adoption of new technologies, although those with large egos may wish to become leaders, testing out the very latest applications and trying to be a state-of-the-art firm. Many of us are fast seconds, if we are lucky. There is a risk in being early for two reasons.

First, you may pick the wrong vendor or wrong technology when several are competing. In the world of Proptech applications, there are several thousand vendors according to Unissu.¹ Unissu is one of several proptech aggregators aimed at matching up investors with proptech startups. CRETech is another,² and Realcomm another, and probably the oldest conduit of information on the integration of technology and real estate aimed at educating real estate operators.³ More recent efforts to match up technology with real estate include Propel by MIPIM.⁴ For any given application there are likely several vendors, and the market will likely whittle these down substantially over time. Picking the ones that will make it is not easy.

- They must have enough capital to last until they reach profitability, something that can take up to ten years in a conservative industry.
- They must have some expertise in real estate, not just great programmers with an idea.
- They must have an application that integrates with other widely used applications like building management systems.
- They must realize that acquiring and filtering data is not easy, and few industry participants want to use multiple vendors, so widespread market coverage is essential.
- Last, they must beat out most of the other vendors, with similar applications as the market has no need for more than a handful of similar vendors.

Second, you might jump in too early while prices are rapidly coming down and by waiting, you could save substantially. For example, the cost of solar photo voltaic (PV) cells has declined, and their efficiency has increased dramatically over the past 45 years increasing on average by about 400 percent.⁵ Stated another way, the cost per watt of electricity produced is about 1/4th what it was four decades ago. Anyone who has installed solar PV cells more than ten years ago, will find they bought units that are now obsolete and worth replacing. The same pattern is true for Light Emitting Diodes (LED) lights. In 2007, a 100-watt incandescent flood light cost perhaps \$6 to \$10 US dollars and lasted 1000 hours at most. An LED bulb of the same lumen output cost around \$89 in 2006 and lasted up to 30,000 hours. Today, that same LED unit will cost less than \$12 US dollars and last up

¹ See PropTech, ReTech, CreTech and Property Ecosystem—Unissu.

² See CRETech.

³ See Realcomm.

⁴ See Propel by MIPIM.

⁵ See M.D. Udayakumar, G. Anushree, J. Sathyaraj, A. Manjunathan, “The impact of advanced technological developments on solar PV value chain”, *Materials Today: Proceedings*, 45, Part 2, 2021, 2053–2058.

to 50,000 hours. When prices were dropping by 20 percent or more per year, it made sense to wait a while rather than be an early purchaser. We can actually model the optimal timing based on the rate of price decline versus the expected savings, so that we maximize our return on investment, but there is no question that it makes sense to wait when prices are declining by 20 percent per year.

Nevertheless, technology will revolutionize all the functions within the real estate industry. Here we provided a brief exploration into some of the applications.

29.1.2 TECHNOLOGY APPLIED TO DESIGN, CONSTRUCTION AND THE METAVERSE

One of the older technology applications, starting in the 1970s, applied to building design is known as **BIM**, for Building Information Modeling. BIM is the digital representation of a three dimensional (3D) physical structure. Within the digital building the level of detail represented can include all of the functional systems in a building such as plumbing, electrical, security, heating, venting, air conditioning, air filters, sensors of all sorts, lighting systems, along with walls, supporting structural components and more. BIM allows designers to understand where there are system conflicts that must be re-engineered. BIM also allows for project management, including planning the staging of equipment and materials on site for the most efficient construction process, or for an interior designer to tweak the lighting systems or a leasing agent to generate a 3D tour of the space.

When BIM is applied to a building that either exists or is in the planning stage, it is known as a “**digital twin**.” When it is not expected to ever be real space, but used only in digital form, it is then simply part of the “**metaverse**.” The Metaverse allows people to play, socialize, collaborate, tweak ideas and plan as if they were together in the same space. It requires high speed and reliable internet access, something becoming more ubiquitous over time. “Meta” means beyond and “verse” is the tail end of the universe suggesting breadth. Some suggest that the term metaverse originated in science fiction, especially Neal Stephenson’s 1992 cyberpunk novel *Snow Crash*, but usage of the term exploded after 2021 when Facebook changed its name to Meta.

Digital-only buildings can be valuable, just like popular domain names can be valuable. The easier they are to find and the more well-known, the higher the traffic on such sites will be. The generation of value is a function of the attractiveness of the digital space to users who wish to tour, visit, interact with each other, sit in on conferences, be entertained and interact with a specific environment. If a digital place provides a user with access to influential people who use the same space, it becomes more valuable. In the entertainment world, think about the value of front row seats at your favorite musical venue. Now, allow several thousand users to sit in that exact seat and interact with the surrounding space in virtual reality. Such a space becomes much more valuable than when constrained to the physical world of one seat one person, and one that requires time to transit to the space, and time to return home. Facilitators of metospace development applications include, as of 2024, Decentraland, Somnium Space, and Otherside, to name a few, with most of these aimed at gaming. Eventually, such sites will enable high quality virtual conferencing and meetings.

Virtual reality experiences require a digital space to be created and can be used purely in a digital environment or combined with an existing physical building. This is called “augmented reality” or **AR** and allows a user to tour an existing real building, or to examine lighting alternatives or different color schemes or different furniture plans, all without leaving home or office. Some people may take little vacations in the future based on augmented reality tours to places otherwise too expensive to visit. This will likely include, via commercialization, all of the world’s major museums and iconic buildings and places.

Among some of the sample BIM vendors as of early 2025 are:

- Autodesk (Revit)
- Oracle Aconex
- Archicad Graphisoft

- Procore
- Houzz Pro

29.1.3 TECHNOLOGY APPLIED TO DATA AND DATA ANALYSIS

It is hard for a student in the current century to imagine the pre-computer and pre-cloud era when data was fragmented, not standardized nor integrated in any searchable fashion. Imagine books and periodicals could only be accessed in person and nothing could be easily copied and shared. Imagine trying to find out what a property sold for in 1975. The analyst would need to start with the address of the property of interest, go down to some form of public record building and look up transfer records, often requiring physically sifting through numerous large-scale books, and then to acquire the physical details would require another set of records kept separately and this process would be done one property at a time.

Any form of analysis, such as mapping or graphs, based on the summarization of data over time or through space began with a very time consuming and costly data collection process. Research and analysis time was dominated by the time required to collect data. Stated another way, 80 percent of the labor required to determine the value of a property or fair market rents was in data collection. Today, current and historical data including property ownership, physical attributes, prior transactions, liens, property taxes, leases, location based social economic and demographic data, image data, flood risk data and more is generally available because the government or private vendors have assembled such data. Private and public vendors have integrated a myriad of data sources that run the gamut from financial to physical to social to legal related information.

Sources of data: We can categorize data as either public or private. Public data is that which is collected by the government, from the local county or city to the province or state to the federal or national. It includes everything we typically find in the census, the Bureau of Labor Statistics, the banking system and more. Some of the public data sources include:

- US Census Data
- Multifamily Data from the US Census
- American Enterprise Collateral Risk Indicators at the Zip Code Level
- Federal Housing Finance Agency data
- CensusHub2

Private data vendors grew as a result of the need to aggregate data from various public sources and over time, more and more private data was added, via direct research and automated spider systems that scan websites and news media. Business, financial, economic, demographic and social data have been combined by Economy.com, now owned by Moody's. In the real estate world we have some data giants like Costar, that have data on all commercial property including apartments spanning leases and asking rents to sales prices and physical attributes, ownership and management contacts and records all integrated with public data on traffic flows and demographics and provided by individual property or in automated graphics making analysis of trends easy. CoStar focused on slightly more institutional property is RCA data from Real Capital Analytics now a division of MSCI. MSCI had bought a similar European commercial real estate data firm, IPD, several years earlier in 2012, while Moody's acquired REIS, another data firm. There are now public data aggregators like CoreLogic® now part of ICE. The reader might notice a trend that larger data firms have been buying smaller data firms and integrating ever more disparate or unique data sets. For example, flood risk data firms (Jupiter Intelligence or First Street as examples) combine flood risk data with fire risk data, weather data and earthquake data to form a more comprehensive and complete solution for users. There are many

other data firms out there in the commercial real estate world and mentioning more of them has risk in that one cannot predict who will merge with whom or be eclipsed by another larger firm. Among these as of 2022, are Compstak/ Realpage, Reconomy, VTS, ATTOM Data, Green Street, and others. In the residential space, there are also firms like Homes.com, Zillow, Trulia, Redfin, Apartments.com as of 2024, that have broadened the geographic scope of market coverage that was once localized, and gained tremendous data utilization efficiencies not possible within only local markets

Last, on the topic of sources of data, we have sensors within and external to buildings that capture images, movement, heat, light, and toxic gases and allow for integrated property management and security management. This real time data must be captured and analyzed and visualized in some sort of building management system, discussed further below.

Standardization and integration: Data comes in different forms, with different errors and with different meta-tags that provide information on when or how it was collected. Among the numerous issues one must deal with when integrating data is to standardize the associated time period of the data. Filtering data so as to eliminate errors is a process all data vendors routinely perform. For example, a lease for \$1000 Euros is recorded by a clerk as \$100. A good filtering system that looks for variance outside of normal expectations may catch the error and allow for correction by generating a red flag for a human to review. In real estate, most data have a geographic component and using a common location identification system (often based on latitude and longitude) is generally a first step, but then such a location identifier is combined with information that may be micro in scope such as a floor, or more macro such as a city, county, state, province, country all of which needs to be attached to the data so that users can sort the data any way they choose. All the major data vendors need to standardize, filter, and attach additional information to each variable of interest to make it user friendly.

Data providers invest considerable programming efforts into automating data analysis. Experts work with programmers to enhance and improve decision driving analysis. For example, graphing sub-lease rates over time might provide a leading indicator of vacancy rates. Vacancy rates over time might provide a leading indicator of asking rents. Providing a tool that allows a user to graph (visualize) all these three together over time and for a specified geographic area, allows market insights that might otherwise be harder to detect. Experts with experience in real estate are required for data vendors to make auto-generated analysis enhancements to databases so that they become more valuable.

An example of data analysis that has evolved over time is the analysis of a trade area and potential demand for a given retail product. Pick any retailer and they will have surveys and estimates on how far people will travel for whatever goods or services they provide. They will assume a certain percentage of patronage based on how they define their trade area. Generally, it will be a decreasing percentage of the market as distance increases from the retail location. Many decades ago, when better household data became available and could be mapped, the trade area was typically defined in distance as the crow flies (straight) without regard to roads. Later on, this was changed into actual travel distance using existing roads, often measured in 1, 3, 5, or even 10 miles or kilometer increments. Then in the early 2000's we changed how we measured distance to time traveled on existing roads, again in increments of 1, 3, 5, 10 minutes and so on, with a decreasing market penetration assumed as the time increased. Note how each iteration of technology allowed measurements to improve and be more reliable. About 2010, some firms started to track cell phone locations and used this data to define trade areas much more accurately. Now, trade areas could be based on data mining, based on tracing the distance traveled for a sample of shoppers, and then matching these households with other socio-economic attributes so that predicting market potential on new sites can be better modeled. Today, cell phone location tracking has become the standard approach to trade area analysis. Someday it will change again or become more refined. The point is that data analysis depends on generating good variables that are related to the prediction or explanation we seek, and that data in general has exploded with all sorts of new sensors being installed everywhere.

Visualizing data via automated graphs or maps has become a huge part of the application tools that users and managers require today. Geographic information systems (GIS) are generally integrated into most applications today along with the data, so that they have become seamless.

Data and analysis proliferation in the real estate market has made the market more efficient. For example, buyers and sellers can find each other faster. All market participants have more information than in earlier decades and the market is more transparent. This has resulted in less market cycle amplitude, faster inventory turnover and more informed forecasting of rents and prices. Better data, real time market and forecasting have dampened the depth of overbuilding or underbuilding. A short caution at the end of this section on the brave new world of real estate data and analysis: The proliferation of new data sets and advanced analytical methods, particularly the rise of Machine Learning (ML) techniques, has triggered an avalanche of new research and insights. A quick Google Scholar search in December 2025 for “Machine Learning” and “Real Estate” yields an astonishing 85,000 results. A resounding success story?

Perhaps, but a closer look through a two-dimensional lens (see Exhibit 29-2) may temper the enthusiasm. ML truly shines when applied to high-dimensional data, abundant in the built environment. The uniqueness and complexity of real estate assets often resist reduction into rigid spreadsheet formats, and the patterns within such data are equally complex—far beyond the reach of traditional linear models like Ordinary Least Squares (OLS). In such cases, the “ML magic” can indeed uncover valuable insights from real-world datasets. So far, so good.

The flipside, however, is that ML’s power diminishes when applied to overly simplified spreadsheet data, where relationships between variables are likely linear and straightforward. In such cases, flashy ML techniques add little value. Don’t be dazzled by fancy terminology; the utility of

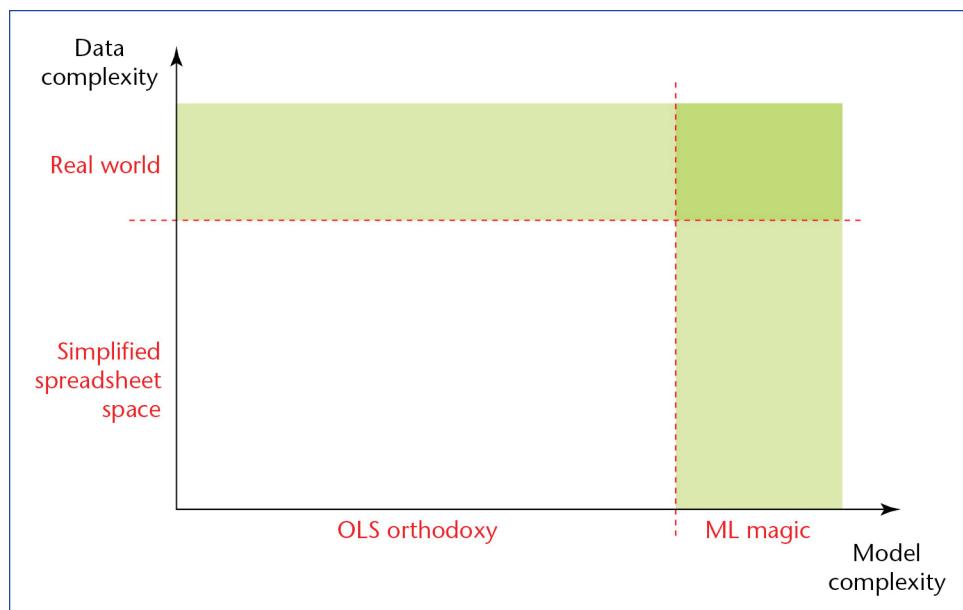


EXHIBIT 29-2 The Sweet Spot of Complex Real World Property Data and Complex Models powered by Machine Learning.

ML depends on the richness and complexity of the data it's applied to and the complexity of the underlying economics.

TEXT BOX 29.1 REAL TIME DATA MINING: THE CASE OF REALPAGE

RealPage, founded in 1998, is a property management software company that offers tools for real estate owners and managers to oversee various aspects of their rental operations, from leasing and accounting to market analytics. Over the years, RealPage has grown significantly, acquiring companies such as Rent Roll, Axiometrics, and Buildium to expand its capabilities. By integrating services and building an advanced platform, RealPage has established itself as a major player in rent setting and property management technology. Its software is used to manage millions of housing units, with its rent-setting tool, YieldStar, drawing both widespread adoption and scrutiny. Journalists have claimed that when so many landlords depend on a single data provider it will lead to higher rents than otherwise. But real time data could also lead to faster rent adjustments in both directions, and there is no empirical proof of price fixing. There are markets where politicians are looking for scape goats for the lack of affordable housing. Most economists will point to supply constraints or strong demand, not to data vendors.

RealPage's rise to prominence can be attributed to its ability to collect, aggregate, and analyze vast amounts of residential rental data. This scale has allowed the company to achieve unparalleled insights into market trends.

As the data set grows, the insights become more precise, and the algorithm becomes more accurate. The use of real time and broad data sets can make markets more efficient, but generally leads to more dynamic pricing. In such informed markets, rents may rise faster and fall faster than in less informed markets. Dynamic real time pricing is already practiced by hotels, airlines, entertainment venues and ride sharing apps, and now has come to the real estate industry. Yet, these other industries have yet to become targets of class action or anti-trust lawsuits.

29.1.4 TECHNOLOGY APPLIED TO MARKETING

Today the buying or selling and leasing of real estate via dominant applications and exchanges has become standard and expected. For residential real estate we think of vendors like Homes.com, Zillow and REALTOR.com, or for residential leasing of apartments, Apartments.com or for house rentals, VRBO and AirBNB, and for commercial real estate we think of vendors like Loopnet or TenX owned by CoStar or Compstak. There are a number of younger firms that have integrated contact management systems with listing platforms, like REALNEX, and virtual touring, see PIX-VR via RealNex, VTS to name a few. Contact management systems, or contact resource management, **CRM**, like Salesforce, are now integrated into other database management systems so that it is easy to connect clients and users to data sharing platforms. As discussed earlier, Virtual Reality, **VR**, spaces are digital representations of 3D space that can be integrated with building designs to allow for self-guided touring and marketing. This is becoming commonplace for both commercial and residential real estate. Again, when VR is combined with real images, like a building, it is called

Augmented Reality, **AR**, a way to overlay what could be by combining digital images with real space images.⁶

Virtual Touring: What all of these successful search engines and touring platforms have in common is near universal coverage of easy-to-use search software that allows a myriad of search filters, a large capital budget allocated to maintaining market share and sufficient capital to buy up competition with innovative and complementary software. Successful firms engaged in the marketing of real estate will provide the user with ever more market information including technology to measure usable space, as a lure to use their site, and these R&D leaders will dominate the market. There is a natural economy of scale for the winners, so we can expect that only a handful of firms will survive for each defined market niche or property type. See Matterport as an example of a service that claims to “Transform any space into a dimensionally-accurate digital twin with your smartphone and the all-new Axis motorized mount.”

Landlords also benefit from new automated services that pull tenant credit and financial information, research the track records of tenants and speed up the tenant screening process.⁷

29.1.5 TECHNOLOGY APPLIED TO OCCUPANCY

Aside from better leasing applications combined with contact resource management, today, there are a host of applications that assist occupants in the more efficient utilization of their space. Since 2010, the space per worker has been coming down steadily with the implementation of unassigned workstations and offices, which are reserved on an as needed basis.⁸ In the US market, this exhibit averaged about 330 square feet per worker from 2000–2008. It surged in 2009–2010 as layoffs lowered the employee count while leases remained in force to 370 and it has been falling ever since. Today many firms use exhibits below 100 square feet per worker on the assumption that “work from home,” WFH, combined with sharing unassigned work stations and offices, will continue to lower the footprint needed to support a workforce that works at home, the office, and everywhere in between. As a result of COVID, the WFH trends accelerated as well as businesses realizing that greater employee productivity and happiness was possible.⁹ For a weekly newsletter aimed at workplace evolution, see IFMA, and Global Workplace Analytics.

Non-real-estate applications like ZOOM, and GotoMeeting, have all facilitated the ability to work from numerous locations. This ability to go into the office less frequently, or not at all, has in turn led to a decrease in the transportation costs per unit of distance, which we have called bid rent curves. The slope of the bid rent curves has flattened and the ability to live further from work has resulted in the growth of tertiary markets, where the cost of living may be lower, or the amenities desired are more affordable. We can expect secondary and third-tier cities to continue to benefit from these trends.

New apps such as those developed by Schneider Electric, or Johnson Controls, monitor every person occupying a building. This enables room heating and air conditioning to be turned off when no one is around, saving energy. Other apps like motion and heat detectors can be used to turn off exhaust fans in underground garages when no one is around, again saving energy.

⁶ Extended reality (**XR**) is a term encompassing VR and AR.

⁷ For example, platforms such as TenantApp.com, Azibo.com, and RentRedi.com illustrate this trend. There are dozens of similar services, though many are unlikely to remain viable in the long run.

⁸ See “Workplace Trends in Office Space: Implications for Future Office Demand” by N Miller in *The Journal of Corporate Real Estate*, 16(3): 2014.

⁹ See “The ROI Of Remote Work” by Kate Lister, 2021.

We have also seen a proliferation of concierge style applications that keep tenants or occupants happy, providing services that range from childcare, dry cleaning, car washing, theater tickets and restaurant reservations, just as higher end hotels might provide. Tenants can now also order, report and schedule repairs or maintenance, provide security information on visitors or deliveries, and all of this automation allows property managers to be more productive and keep tenants happy. See for example HqO.

Last, we have a plethora of co-working applications that commit to longer-term master leases and sublease to smaller tenants for a shorter period of time, even days or hours. These spaces cater to the tenant that needs overflow space or space in a variety of locations. Some add social aspects to the space and provide extra shared services and amenities. An example of some of these applications includes those who have been around for a while like Regus, Liquid Space and WeWork as well as Knotel and many others.

29.1.6 TECHNOLOGY APPLIED TO BUILDING MANAGEMENT SYSTEMS AND OPERATIONS

Facilities management has benefited as much as any other aspect of real estate, lowering costs and improving services. It may help to start with an example. Historically, elevators were placed at pre-determined locations where they waited until someone pressed a button. A 20-story office building or apartment building may have had, as the default, one elevator stationed at floor 1 or zero in Europe, one elevator at floor 10, and one at floor 20. Whenever no one pressed a button to summon the elevator for more than 60 seconds, these three elevators would return to pre-determined locations. Such was the state of the art a few decades ago. Now, by collecting the data on which floors receive the most requests monitored by minute and hour over the course of several days and continuously updated with a tolling database, a frequency table and pattern can be extrapolated. This frequency table is then used to determine where the elevators will wait at any particular time of the day and based on that particular weekday. Instead of returning to the old, scattered pattern, now the elevators default to floors more likely to need elevator transit. The result is less waiting time for the building occupants and less energy used as now the elevators travel less distance than before. It is a win-win for occupants, managers, and owners. Such software is one of many operational apps that allow a building management system, **BMS**, to operate more efficiently and with better service. Now imagine the same type of analysis applied to temperature zones in a building with automated shading and venting depending on outside and inside temperatures, or filters that measure various gases and or particulates. Lighting, shades, elevators, work orders, filters, security systems will all operate better with numerous sensors for images from cameras and biometric readers, automated human temperature readers, heat, carbon monoxide, dust sensors, air pressure sensors, energy consumption, water consumption, and so forth. One can attach a sensor to every fire extinguisher eliminating the need for inspections, or to every light to let maintenance know if the light is about to burn out. The cost of such sensors is modest compared to the benefits of automating comfort and security and safety within a building.

Among the systems that many BMS will integrate into one monitoring system include, but are not limited to the following:

- HVAC
- Lighting control systems
- Water meters
- Gas meters
- Power meters
- Energy meters
- Solar energy systems
- Weather station
- Benchmark databases

- Elevators and escalators
- Smoke curtains & valves
- Fire extinguisher
- Fire detection
- Intrusion detection
- Access control
- CCTV (Cameras)
- Motion Detectors
- Leak Sensors
- Maintenance requests
- Work order status

BMS is available from a variety of vendors and the most successful firms will have systems that are scalable, easy to add apps, and well-integrated with systems that are easy to manage. While the vendors landscape will be in flux for a while a small sample of vendors are provided below. Yardi is among the oldest of the firms below and most established.

29.1.7 TECHNOLOGY APPLIED TO MEASUREMENT AND ANALYSIS

Geography: Real estate not only has a global identifier based on latitude and longitude, which is the standard requirement for geographic information systems, GIS, and mapping systems, but also a whole host of other identifiers. This can include political geographies, city, county, province, state, metro, country, and school district. There are also geographic spatial estimates such as 100-year flood plains, or fire risk indices that are tied to geographic space. All of these allow for the generation of metrics that feed into analytical decision models or reports. GIS has become so integrated into analytical platforms that we barely notice it anymore.

Space Layout and Size: More recently, new technology software applications allow for the measurement and layout of buildings. Cameras (or phone cameras) and lasers can be used to judge distance, tell us the width of a room or height of a wall, and even draw layouts. This allows a user to estimate more accurately the actual square footage or square meters available for occupants. Matterport was mentioned above in the discussion of virtual touring, but such applications also allow for measurement of space and reproduction of 3D models in digital form all based on a series of images captured via a smart cell phone.

Sensors and cameras: Sensors and cameras have been previously mentioned as allowing for motion detection, gas detection, pressure detection, air quality monitoring or security investigation and management. Critical to the effective use of sensors, such as those that measure temperature, or energy consumption, are visualization tools and exception reporting based on an analysis of what is normal and what is unusual. Data from sensors or cameras would quickly become overwhelming for a human to monitor, so most software aggregates information and reports it in lumps, like energy

EXHIBIT 29-3 Building Management System Vendors (a Small Sample).

Firm	Website
• JLL Building Engines	www.buildingengines.com
• Appfolio	www.appfolioinc.com/
• RealPage	www.realpage.com/
• Lessen	www.lessen.com
• VTS	www.vts.com
• Yardi	www.yardi.com/

consumption over some period of time and compares it to benchmarks based on past data. Such software is typically set up for notification and visualization, not unlike a Ring security system, but with more diagnostic screens so that a variety of comparisons to normal can be considered.

One example of data aggregation and reporting are those applied to energy consumption, water consumption and carbon footprint estimates. For example, most solar photovoltaic systems like Enphase will report on the production and consumption of electricity over the course of a day, week, month or year. This allows the user to know if they have enough solar cells to cover their needs or to judge the impact of energy savings techniques applied to certain times of the day, known as demand reduction or shifting. Another example aimed more at commercial real estate, would be Measurabl, that helps connect data from gas, electric, water or other meters in such a way as to provide automated reports and estimates of carbon impacts, water and energy use. Measurabl allows environmental and energy related reporting. Another similar firm is Verdani that assists in all sorts of environmental reports.

Prior to automating the measurement systems, such reporting was difficult or impossible. To this date, many occupants in buildings have no submeters for gas, electric or water consumption and this lack of matching consumption with user costs results in a lack of incentives for savings. Those who are less efficient are subsidized by those who are more efficient, and the only way around this is to be able to measure individual occupant consumption. There are countless other examples of measurement and reporting applications, but in most cases, time savings, ESG goals and better management results are the driving forces.

29.1.8 TECHNOLOGY APPLIED TO OWNERSHIP FORMS

Individuals have long broken up real estate interests into formats that allow multiple owners. Often this has been via a syndication of some form that creates a partnership with active general partners and passive limited partners. Such structures may allow for a few to a few hundred partners. When the entity seeks greater investor participation, two other formats are often used. One is a real estate operating company, REOC, essentially a corporate form of ownership applied to real estate. An example of a REOC would be the Marriott hotel chain. The other format, with tax advantages over a REOC is a REIT, real estate investment trust. REITs were authorized in 1960 and became legal in 1961 in the US. Since then, they have grown and become global in scope¹⁰ with REIT formats now allowed in numerous countries. Their advantage is the pass-through taxation so that there is no double taxation as with a corporation. In order to secure this tax-benefit, the REITs must pass through the majority (90 percent) of their taxable income, and this limits the ability to retain earnings.¹¹

In more recent years, there have been two newer forms of ownership aimed at greater democratization of real estate investing or at least greater efficiency in the management and transfer of ownership interests. Think of democratization as allowing smaller investors to get into real estate, similar to the original goals of REITs when first formed. These two newer forms are **crowdfunding** and **tokenization**. Each will be discussed in turn.

Crowdfunding: Crowdfunding was effectively established as a legal format to own businesses in 2013. The Jumpstart Our Business Startups Act, or **JOBS Act**, is a law intended to encourage funding of small businesses in the United States by easing many of the country's securities regulations. It passed in the US on April 5 of 2012. Title III, also known as the CROWDFUND Act, legal as of May 16, 2016, created a way for companies to use crowdfunding by issuing securities.¹² The impact

¹⁰ Public EQUITY REITs were about \$1.75 Trillion US Dollars in scope as of 2024.

¹¹ See www.REIT.com

¹² About \$20 to \$30 Billion has been raised through crowdfunding as of 2022.

EXHIBIT 29-4 Sample Crowd funding platforms (not an endorsement)

• Crowdstreet	www.crowdstreet.com
• YieldStreet	www.yieldstreet.com
• Roofstock	www.roofstock.com
• Fundrise	www.fundrise.com
• Cadre	www.cadre.com
• Realtymogul	www.realtymogul.com
• Realcrowd	www.realcrowd.com

of the JOBS Act has been a plethora of fundraising aimed at smaller investors. Among the several crowdfunding platforms, two of the more successful platforms have been Realty Mogul and Crowdstreet. As of the end of 2024, Realty Mogul had raised nearly one billion US dollars, invested over one billion dollars in property, with 250,000 plus members. This suggests an average investment of just under \$4,000 US dollars. By late 2024, Crowdstreet had raised \$4.4 Billion US dollars and bought 790 properties.

Tokenization: Tokenization is the most recent of the ownership forms. “The success of cryptocurrency and blockchain technology undergirds asset tokenization. Real estate tokenization is a natural evolution of decentralized finance.”¹³ Traditional commercial real estate investing including direct investing and crowdfunding has been fairly illiquid. Liquidity requires a well-informed secondary market. By improving the efficiency of recording and transferring ownership interests, leases, financing and all property records, one might lower transaction and asset management costs. Relying on blockchain for verification of such interests, secondary market liquidity may become possible. That is the central benefit of tokenized real estate, the possibility of lowering asset management costs, transaction costs, and the ability to dissect an interest into any sized slice that suits the purpose of investors and transfer it without extensive government recording.¹⁴ In a sense, **if tokenization succeeds, it will be partly because of the disintermediation of the government as a sacred place for record keeping.** An entity that owns a property will still be recorded via government recording procedures, but the interests in this entity might be traded without government recording as long as the owners feel their interest is safe via a blockchain distributed ledger system.

To tokenize a property is to break it into small pieces of ownership. For example, a 100,000 square foot building worth \$10 million might be broken into 100,000 pieces where each one has a base value of \$100. This makes the asset more accessible to smaller investors, a process known as the democratization of real estate. Of course, the asset in tokenized form may trade for more or less than the asset if sold as a unit. Just as the capitalized market value of a REIT may differ from the underlying asset values, so too will tokenized values vary over time. One should also note that instead of trading in dollars, the currency will more likely be in Ethereum, (such as ERC20, a security token), which makes recording, tracking and trading ownership shares much easier and more efficient. Secondary exchanges where such tokenized interests can be traded will eventually come into being and the current stock exchanges are investing in the technology to facilitate tokenized trading, not just in real estate but all sorts of assets.

Real estate tokens are digitalized economic rights to the assets as listed on the blockchain.¹⁵ The assets themselves and their corresponding legal constructs are off-chain. Legal linkage and

¹³ Guy Tcheau, Principal Investors, May, 2022 presentation on Tokenization.

¹⁴ Whether securitized tokens could be used in 1031 exchanges remains to be seen.

¹⁵ Keep in mind that blockchain is a cloud based immutable recording of interests. Blockchain is essentially a distributed data structure that holds transactional records in the form of a chain of records (blocks) that are linked to one another. The contents of each block and the entire chain is secured by a cryptographic signature known as a hash.

recourse between on-chain digital tokens and their off-chain real world asset counterparts are still evolving and uncharted. Real estate tokens are deemed securities and must comply with securities regulations. Regulatory and compliance layers must be incorporated into tokenization process. In theory, tokens could be used to fund debt, equity or different slices and interests in real estate. For example, slices of CMBS could be tokenized, including the residual piece.

A token's interests will be defined via a smart contract. Tokens could be listed on exchanges and various investors would be able to continually reprice the value of the interests. Such a secondary market would enhance market efficiency and reduce market friction, as the costs to trade would be far less than typical in today's real estate market. Over time, vendors will supply market information that helps investors price tokens and determine whether to sell, hold or buy. Properties that have been tokenized may switch managers and leasing staff but might retain the tokenized format for long periods of time.

Countries leading in tokenization are Germany, Singapore, Israel, Switzerland, Lithuania and Liechtenstein, as of 2022.¹⁶ Within the US, the states of Wyoming, Texas are leaning towards legalizing blockchain supported forms of ownership. Tax havens may soon follow with tokenization strategies like Malta, Estonia, the Cayman Islands, and Gibraltar. One of the first completed US based tokenization deals was Aspencoin. It is an LLC token-based ownership of the St. Regis Aspen Resort in Colorado. 18 million tokens were sold at 1.00 US Dollars and the minimum investment was 10,000 US Dollars accepted from qualified accredited investors.¹⁷ Ethereum was used for the blockchain platform.

One example of a US based tokenization platform is Digishares. Many others will follow.

29.1.9 TECHNOLOGY APPLIED TO TRANSACTION MANAGEMENT AND REPORTING

In the past, real estate closings were mammoth events filled with attorneys, closing officers, buyers and sellers and brokers who met and signed in person. When someone needed to sign remotely, notaries were relied upon to verify the signer's identity. Accelerated by COVID in 2020–2022 DocuSign came along and helped to automate remote signatures. Momentum and a lack of questioning the pre-existing process resulted in a huge waste of time for many parties. Generally, there is no need for real estate agents to be involved in a closing, other than to pick up a check which can now be electronically deposited. Buyers and sellers could sign without being in the room at the same time. Closing statements and documents could be prepared and reviewed and approved ahead of time. We can expect remote closings in the future, managed by title companies, lenders and brokers.

Prior to closings, major time impediments are the time required for traditional appraisal reports, title search and closing document preparations. Today, there are firms that have automated all of these processes. Not all transactions can be accelerated, but many can be accelerated. Appraisals can be done with automated appraisal models known as AVMs or CAVMs for Automated Valuation Models or Commercial Automated Valuation Models, as long as the property information is available, and the property is not unique in the factors driving its value. While due diligence generally requires personal inspections, some of the process can be done remotely with virtual touring, images and certifications by facility managers on the condition of building systems. Drone images and video can be used to inspect roofs and maintenance records can be captured and shared. The point is that closings in the future will occur faster and more efficiently than in the past.

¹⁶ Guy Tcheau, Principal Investors, and authority on tokenization.

¹⁷ Accredited investors are determined by net worth and income. Generally accredited investors must have net worth of \$1 million US dollars excluding the personal residence and \$200,000 plus of income per year.

EXHIBIT 29-5

Selected AVM Providers (not an endorsement)

Firm Name	Property Sector	Website
Altus	Commercial	www.altusgroup.com
Attom Data	Residential	www.attomdata.com
Collateral Analytics (Now ICE)	Commercial and Residential	www.icemortgagetechnology.com/solutions/valuations
CoreLogic	Residential	www.corelogic.com
Green Street	Commercial	www.greenstreet.com/
Real Capital Analytics now MSCI	Commercial	www.msci.com/our-solutions/real-assets/real-capital-analytics
RedIQ	Commercial	www.rediq.com/
REIS/Moody's	Commercial	www.moodyscrc.com/
REONOMY	Commercial	www.reonomy.com/
Rockport VAL	Commercial	www.RockportVAL.com
Valcre	Commercial	www.valcre.com
Veros	Residential	www.veros.com

Reporting is a huge expense for all asset managers. A popular software app for bookkeeping is called Quickbooks. Quickbooks feeds into another popular software app called TurboTax. In the real estate industry, the same type of parallel reporting systems has been developed for property to portfolio level reporting, which then feeds into investor or owner level reporting. Examples include Appfolio and Juniper Square, both of which combine some property management accounting with investor reporting. Such applications save accounting time and expedite reports, allowing asset managers to communicate more efficiently with investors.

Such reporting automation platforms save enormous administrative and communication time, allowing owner/managers to consider more partners and to keep passive investors well informed. These platforms also speed up tax related reporting, such as K-1's and allow passive investors to check for updates with personalized platforms.

29.1.10 THE POWER OF MULTIPLE GRADUAL IMPROVEMENTS IN PROPTech AND EASE OF INTEGRATION

The transformative potential of Proptech often lies not in singular groundbreaking innovations, but in the cumulative impact of multiple gradual improvements. These incremental changes, when executed effectively, reduce transaction costs, streamline processes, and create entirely new markets. The whole is more than the sum of its parts. Integration of new proptech applications will be most successful if they integrate with existing widely adapted software like MRI or Yardi as examples.

Airbnb is a prime example, becoming one of the dominant short-term rental platforms, along with VRBO, by addressing multiple facets of short-term letting—that is, advertising, screening, contracting, payments, insurance—none of which were revolutionary on their own. These solutions were adapted from other e-commerce sectors, yet their combination reduced friction in the rental process so effectively that Airbnb created a multi-billion-dollar industry virtually from scratch.

29.1.11 THE MERITS OF MIXED-SKILLS TEAMS IN DATA-RICH REAL ESTATE INVESTMENTS

Evidence-based decision-making in data-rich real estate investment situations requires teams that combine diverse skill sets. Advanced empirical capabilities must be matched with real estate expertise to produce insights that are both robust and actionable.

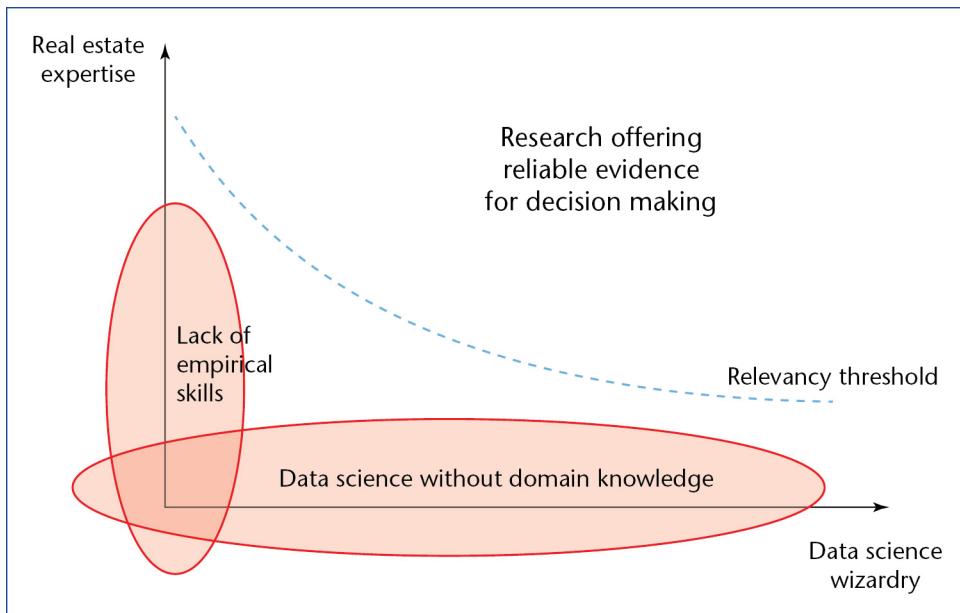


EXHIBIT 29-6 The Combination of Real Estate Expertise and Empirical Skills Needed for Evidence Base Decision Making.

On one hand, domain experts bring critical understanding of the economics, market structures, legal frameworks, and regulatory environments that define real estate. They “get real estate” and understand the practical realities of the industry, from the intricacies of the investment process to the nuanced expectations of developers, tenants, and investors. However, a team composed solely of real estate veterans may (with all respect) lack the empirical and technical skills required to harness the full potential of modern data analytics. While advanced algorithms and user-friendly software have become more accessible, applying them effectively requires a deep understanding of the principles behind them. Without this foundation, amateurish misuse can lead to costly errors—“A fool with a tool is still a fool.”

On the other hand, a team composed entirely of data scientists may excel at handling large datasets, applying sophisticated algorithms, and uncovering patterns. However, without a strong grasp of the realities of real estate, they risk asking the wrong questions, misinterpreting results, or missing critical context. They may not fully appreciate what makes or breaks a good investment, what stakeholders value, or the practical and legal constraints of the market. Ignoring these factors can render even the most technically sound analysis irrelevant or misleading.

A world with more and more Proptech requires **mixed teams** that blend complementary skills. Domain experts ensure the relevance and practicality of analyses, grounding the work in real-world knowledge. Data scientists contribute technical expertise, enabling the use of advanced methods and rigorous data analysis. Together, such teams foster a collaborative approach that bridges the gap between insight and application, ensuring that evidence-based decisions are both technically robust and practically viable.

By leveraging the strengths of both groups, mixed teams can maximize the potential of data-driven insights, producing outcomes that resonate with the complexity and demands of the real estate sector.

29.2 CHAPTER SUMMARY

As of 2025, there are several thousand proptech firms collecting data, processing information, providing software as a service, and helping to make the real estate industry more efficient. Many of

these firms will not be in business in five years, but those that survive will change the industry for the better. Automated buildings are healthier, safer, more productive, touchless, and more efficient to operate. Machine Learning, cameras, drones and sensors all add internal data to the data provided by financial and accounting systems or brought in from external vendors on the market and the economy. Drones, cell phone tracking, foot traffic, visual images, VR, facial recognition and expert systems can improve actionable decisions on such processes as marketing, occupant productivity, security, realistic trade areas, or internal energy optimization.

As mentioned above, there are many firms out there trying to help with the digitization and automation of real estate. At a minimum they must be secure and work with existing platforms, especially building management systems or existing accounting systems. ESG and sustainability are new reporting priorities and requirements and some of these go hand in hand with wellness and occupant productivity which are enhanced by these Proptech firms.

KEY TERMS

- AR
- AVM
- BIM
- Blockchain
- BMS
- CAVM
- CRM
- Crowdfunding
- Digital Twin
- Metaverse
- Proptech
- Tokenization
- VR

STUDY QUESTIONS

Conceptual Questions

29.1 For each of the following functional areas, find at least one proptech firm and answer the questions below in the columns.

Functional Area	Firm Name	Year Established	Market Coverage	Describe Product
Data				
Leasing				
Selling Real Estate				
Designing & Building				
Valuing				
BMS				
Climatetech ESG				
Tenant Experience				

- 29.2 Describe sources of data that might feed into a property management system in the future?
- 29.3 How has retail trade area analysis changed over time with new and better technology?
- 29.4 What is a building management system? Describe how it works? Find a vendor that has a BMS. Does it include security and access control? Do you think that will become more important in the future?
- 29.5 Describe how COVID from 2020 to 2022 has impacted the use of space by office tenants? What are examples of environments or features that are important to tenants? What kinds of jobs require in person collaboration and what types of jobs can be done remotely? Do you think office space utilization will increase over time?
- 29.6 What is a digital twin? How does this allow for better leasing, or building management?
- 29.7 Find an example of a digital or virtual only platform. How can a virtual only place have value?
- 29.8 Go to a proptech aggregator such as Realcomm, CRETech, MIPIM Propel, or MetaProp and describe the topics covered in their next or most recent conference or newsletter.
- 29.9 What legal change or act in the US allowed for crowdfunding to flourish?
- 29.10 See if you can find a platform that facilitates tokenized real estate investment?

Appendix: Additional well established data vendors:

See ESRI at GIS Mapping Software, Location Intelligence & Spatial Analytics | Esri one of the leaders in mapping and spatial data analysis firms.

For a product geared more towards retail see the Site to Do Business, STDB, which works with ESRI data at Site To Do Business: for Commercial Real State Professionals (stdb.com) There are simpler products like Maptitude built to work with Excel, where you import your own data, see Maptitude GIS and Mapping Software (caliper.com) or Tableau at Business Intelligence and Analytics Software (tableau.com). There are also vendors who do nothing but integrate disparate sources of data, aimed at serving the real estate industry like Cherre, see Cherre, a software application that assists in the process of standardizing and integrating useful data.

ADDITIONAL REFERENCES RELATED TO CROWDFUNDING VIA TOKENIZATION

Best Real Estate Tokenization – SolidBlock www.crowdfundinsider.com/2021/04/174078-istox-tokenizes-another-fund-mapletree-merit-fund/www.techinasia.com/kkr-eyes-4m-tech-fund-tokenization-sg-platformwww.finews.asia/finance/36863-temasek-s-fullerton-to-tokenize-private-equity-fund-on-addx
www.alternativeswatch.com/2021/10/01/digital-exchange-addx-tokenizes-partners-group-private-equity-fund/
<https://forkast.news/headlines/addx-investcorp-tokenize-us-real-estate/>