

Chapter 16 Supplement

Some background about investment risk, return, the SML, and the OCC for discounting...

David Geltner

Four cases of risk and return:

When outcomes are subject to a truncated probability distribution, as with options and flexibility, then the OCC (which is the discount rate determining the present value) typically varies across “states of the world”. A “state of the world” is a node in the binomial lattice, a particular combination of a point in time, relative to the expiration of the option, and a current value or price of the underlying asset or “state variable”. (The “state variable” is the value of the real asset that can be obtained or given up or delayed by exercise of the option, or the value of the commodity that can be produced, such as copper in the case of a copper mine.) So, we don’t use risk-adjusted discount-rate discounting to determine present value (PV); we use certainty-equivalence discounting instead, as described in section 16.3.2 of Chapter 16. And then if we want to know the OCC, we “back out” the OCC applicable to each state of the world AFTER we have determined the value of the option in that state.¹ The OCC is just the ratio of the expectation of the future value of the option next period divided by the PV of the option in the current period and state (minus one):

$$1 + \text{OCC} = E[\text{OptionValNextPeriod}] / \text{PV}[\text{OptionInCurrentState}].$$

Let’s label “ $E[\text{OptionValNextPeriod}]$ ” as “X”, and let’s label “ $\text{PV}[\text{OptionInCurrentState}]$ ” simply as “PV”. In terms of signs (positive or negative), there are four possibilities for the OCC and X:

- **Case 1:** X & OCC both positive
- **Case 2:** X positive, OCC negative
- **Case 3:** X negative, OCC positive

¹ There are some types of claims for which literally there is no OCC in the form of a risk-adjusted discount rate. Consider typical futures or forward contracts. These are structured such that no cash changes hands up front, hence, $\text{PV}=0$. Yet these contracts typically have non-zero expectations on their future cash flow. No discount rate can mathematically discount a non-zero future expectation to a zero present value. Such contracts must be valued using certainty-equivalence discounting.

- **Case 4:** X & OCC both negative.

If the OCC is positive, then let's say we call that "positive risk" (ignoring for ease of illustration the rare case where the OCC is positive but less than the riskfree interest rate). If the OCC is negative then we have "negative risk". Consider each of the four cases...

Case 1. If the expected payoff of the asset is positive ($X > 0$) then "positive risk" ($OCC > 0$) corresponds to the situation where the investor is worse off the greater the risk (larger OCC), hence the investor won't pay as much for an asset that has a positive expected value next period but more positive risk:

$$PV = X/(1+OCC) < X.$$

$$\text{e.g.: } 100/(1+10\%) = 100/1.1 = 91 < 100.$$

A typical Case 1 situation is where the asset being valued has a range of possible positive outcomes that co-vary positively with the investor's wealth. Thus, the asset tends to pay off well when the investor's other wealth is also paying off well, and the asset tends to pay off worse (lower) when the investor's other wealth is also paying off poorly. Therefore, the asset adds to the volatility or variability (or uncertainty) in the investor's future wealth, and this is a detraction because the investor is presumably risk averse (doesn't like volatility in her future wealth). This is why the investor doesn't want to pay as much for assets with the same outcome expectation but more positive risk (range of outcomes), producing the lower PV. This situation, positive X and positive OCC, is the most common situation, and applies to almost all "underlying assets", such as stocks, bonds, real assets, most commodities.

Case 2. If the expected payoff of the asset is positive ($X > 0$) then "negative risk" ($OCC < 0$) corresponds to the situation where the investor is better off the greater the risk (larger absolute value of OCC), hence the investor will pay more for an asset that has a positive expected value next period but more negative risk:

$$PV = X/(1+OCC) > X, \text{ since } X > 0 \text{ and } OCC < 0.$$

$$\text{e.g.: } 100/(1-10\%) = 100/0.9 = 111 > 100.$$

Case 2 situations are more rare or special. They involve the value of the asset co-varying negatively with the investor's wealth. A typical example would be a put option on an asset with positive risk, such as a put option on a typical stock or bond or real asset. Such an option allows the asset to be sold for a specified value. Thus, as the value of the underlying asset falls, the value of the put option rises. Typically most of an investor's wealth will be in long positions in assets with positive risk, such as stocks, bonds, or real assets. If the investor holds put options on those assets, then as the value of his wealth falls due to drops in the prices of his assets, the value of his put options rises, partly offsetting the drop in his wealth value, thereby reducing the variability or uncertainty in his future wealth value, providing a type of "insurance" for his wealth. The "negative risk" in the put option is therefore a "good thing" for the risk-averse investor, and he will pay more for an asset with positive expected payoff and more negative risk (holding other things constant). In general, insurance policies are examples of Case 2 situations where X is positive and OCC is negative. And Case 2 situations usually have to be contrived or constructed by means of some sort of financial engineering or insurance structuring.

Case 3. If the expected payoff of the asset is negative ($X < 0$) then "positive risk" ($OCC > 0$) corresponds to the situation where the investor is better off the greater the risk (larger OCC). Negative-payoff situations generally refer to future costs or liabilities. Future costs or liabilities have negative present values, and therefore an investor must be paid to take on such future obligations (e.g., selling short, or taking out a loan). Investors require NPV at least equal to zero, so a negative PV would require that the investor must be paid a positive amount at least equal to the absolute value of the negative PV of the future obligation in order to make the investor willing (or indifferent) to take on the obligation. The lower the (positive) OCC, the more you have to pay the investor to take on the obligation, whose (negative) outcome is more certain with lower OCC (that is, with lower risk):

$$PV = X/(1+OCC) > X, \text{ since } X < 0, \text{ and } OCC > 0.$$

$$\text{e.g.: } -100/(1+10\%) = -100/1.1 = -91 > -100, \text{ but:}$$

$$-100/(1+5\%) = -100/1.05 = -95 < -91 \text{ (and } +95 > +91).$$

A typical case is construction costs in a real estate development project. If the costs are more certain, then they have a greater negative present magnitude and bring down the NPV of the

project. On the other hand, suppose the construction costs co-vary positively with the real estate value (positive risk, hence $OCC > 0$). Then, when the real estate being built does well, the construction costs will also tend to be higher, holding the increase in profit down; but conversely if the real estate being built does poorly (low RE value) then the construction cost will also tend to turn out lower than previously expected (positive risk, $OCC > 0$) and this will tend to preserve the profitability of the project. Thus, greater positive risk (greater OCC) in the construction costs corresponds to less risk in the construction project's outcome. And this equates to a higher NPV for the project *ex ante* (other things being equal, in particular, for same given expected construction cost of amount "X"). For example, suppose the PV of the real estate to be built is \$120, and the expected construction cost is $X = -100$. Then if the construction cost is riskless and the riskfree rate is $r_f = 3\%$, then the NPV is $120 - 100/1.03 = 120 - 97 = \23 . But if the construction cost has risk warranting $OCC = 10\%$ then the NPV is $120 - 100/1.1 = 120 - 91 = \29 .

Case 4. If the expected payoff of the asset is negative ($X < 0$) then "negative risk" ($OCC < 0$) corresponds to the situation where the investor is worse off the greater the risk (larger absolute value of OCC). As in Case 3, the negative-payoff situation refers to future costs or liabilities, hence, a negative PV for which normally the investor would have to be paid a positive value at least that large in order to make the investor willing (or indifferent) to take on the obligation. The more negative (larger absolute value) is the (negative) OCC , the more you have to pay the investor to take on the obligation, whose (negative) outcome is less certain (with more negative OCC , that is, more negative risk):

$$\begin{aligned} PV &= X/(1+OCC) < X, \text{ since } X < 0, \text{ and } OCC < 0. \\ \text{e.g.: } &-100/(1-10\%) = -100/0.9 = -111 < -100, \text{ but:} \\ &-100/(1-5\%) = -100/.95 = -105 > -111 \text{ (and } +105 < +111). \end{aligned}$$

Case 4 situations are not typical when we're dealing with underlying assets or with options on assets, where we typically assume that there is complete "limited liability" for the investor, that is, the investor will never experience an outcome worse than zero. (e.g., You buy a call option on a stock; the stock may become worthless or your option may expire out of the money, so you may end up with zero, but that's the worst you can do; you won't actually have to pay out more cash.) However, if we're dealing with options on operating cash flows (instead of on assets), then Case 4 situations can arise, because the operational constraints may force the investor to experience

negative payoffs, and those payoff possibilities may be uncertain in a manner that co-varies positively with the investors wealth thereby exposing the investor to greater risk. The negative OCC is then due to the negative nature of the outcome possibilities and the fact that the greater absolute risk is a “bad” thing for the investor.²

Consider a copper mine example. Suppose that copper has “positive risk” and that copper prices tend to co-vary positively with the investor’s wealth, such that copper prices tend to be high when the investor’s wealth is otherwise high and copper prices tend to be low when the investor’s wealth is otherwise low. The copper mine has a delay or “pause-production” option reflecting the flexibility to stop production if/when copper prices fall too low. This saves money compared to if the flexibility did not exist (because you save variable costs), but the mine is still subject to certain fixed costs which cannot be avoided. If copper prices are sufficiently high then the expected outcome of the mine operation next year will be positive and we’re not in a Case 4 situation. If copper prices are sufficiently low such that there is no chance production would be resumed next year then we’re in a Case 3 situation of negative expected outcome (the fixed costs) but no uncertainty, hence, $OCC = r_f > 0$. But there can be intermediate situations in which the copper price is low enough such that the expected outcome next year is negative, yet there is uncertainty because the price may be sufficiently high so that the mine will lose less money by going into operation to at least partly defray the fixed costs. In this situation, the greater the range in outcomes, the worse off the risk-averse investor is today, holding all else constant (in particular, holding the expected outcome, “X”, constant).

To see this, consider two different mines, or two different mine production protocols, subject to two different levels of fixed costs and two different rates of copper production and/or of variable costs providing different profitability functions. In both cases the expected outcome next year is -100 (given the low copper price in the current state). In one case the range is +/- 10, between fixed costs of -110 on the downside or net operating profit of -90 on the upside, that is, a loss of 90 on the upside. (And we’re assuming equal chance of either up or down.) In the other case the range is +/- 20, between fixed costs of -120 and net operating profit of -80 corresponding to a loss of “only” 80. (Of course, we’re in a state of the world where copper prices are pretty low, which is


² Case 4 situations can also arise with some positions in some derivatives contracts, such as short positions on positive-risk underlying assets that have moved in an unfavorable direction since the contract was written.

why the future is looking so bleak.) The downside outcome occurs when copper prices are down (further), which since copper co-varies positively with the investor's wealth is also when the investor's wealth is likely to be down. This also would be the downside outcome for the option, resulting in the -110 outcome for the mine (or protocol) with less negative risk (less absolute value OCC), but an even worse outcome of -120 for the mine (or protocol) with the greater negative risk (OCC that is "more negative", that is, a bigger negative number). Just the opposite occurs in the upside outcome, where a higher copper price, corresponding to a higher likely investor wealth, is also associated with the less negative outcome in the mine (-90 for the less negative risk mine; -80 for the more negative risk mine). Thus, we see how the "negative risk" in this case is associated with positive covariance between the mine outcomes and the investor's wealth, and how the mine with the greater absolute value of range of outcomes increases the investor's wealth volatility more than the mine with the lesser absolute value of range of outcomes:

	Upside Outcome	Downside Outcome
Copper Price	Low but not as low	Even lower
Investor's Wealth	Higher	Lower
Less Negative Risk Mine	-90	-110
More Negative Risk Mine	-80	-120

Note that in this Case 4 situation, even though the option's outcome co-varies positively with the underlying asset and with the investor's wealth, the OCC is negative, implying "negative risk" the way we have defined this term (though this is only a labeling convention). This negativity is due fundamentally to the negative expectation in the option outcome ($X < 0$), which in turn is due to the lack of complete limited liability in the investor's exposure. Here we see how this can happen in real options analysis modeling operations rather than assets. It can also happen in financial engineering with certain situations in hedges using derivatives.

Concave (declining marginal utility of wealth) Utility Function (aka “risk aversion”) Implies that Risky Negative Future Values Are Worse (lower expected utility) Than Riskless Negative Values...



Note: This is an example of “Case 4” in the table on the following slide...

Note: Utility function is equivalent to any affine transformation of itself such that: $U^* = a + bU$, where “a” is any value and “b” is any positive value; and above result still holds for U^* .

Opportunity Cost of Capital (OCC) as Risk Adjusted Discount Rate (RADR) Relating Future Expected Outcomes to Present Value (PV) with Risk-Averse Utility (concave utility function: declining marginal utility of wealth) or Equivalent Mkt Pricing (SML)...

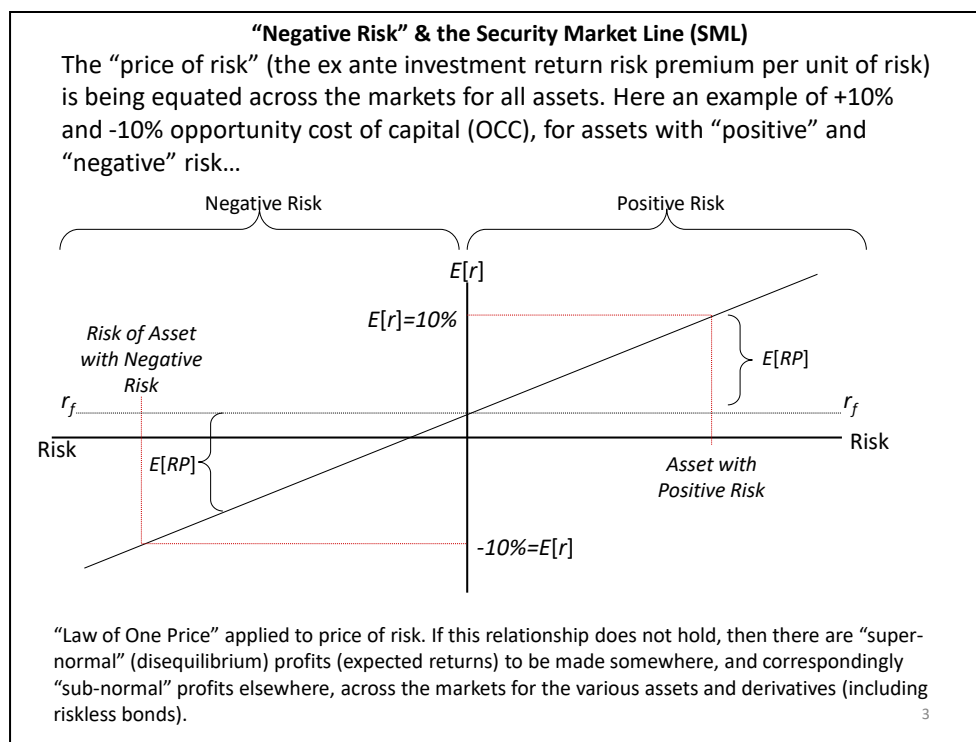
Case 1: Positive Expected Outcome, Positive OCC	Case 2: Positive Expected Outcome, Negative OCC	Case 3: Negative Expected Outcome, Positive OCC	Case 4: Negative Expected Outcome, Negative OCC
$E[\text{Outcome}] \text{ “X”} > 0$ OCC “r” > 0 $PV = X/(1+r) < X$	$E[\text{Outcome}] \text{ “X”} > 0$ OCC “r” < 0 $PV = X/(1+r) > X$	$E[\text{Outcome}] \text{ “X”} < 0$ OCC “r” > 0 $PV = X/(1+r) > X$	$E[\text{Outcome}] \text{ “X”} < 0$ OCC “r” < 0 $PV = X/(1+r) < X$
X=100, r=10%: $PV = 100/1.1 = 91 < 100$	X=100, r=-10%: $PV = 100/.9 = 111 > 100$	X=-100, r=10%: $PV = -100/1.1 = -91 > -100$	X=-100, r=-10%: $PV = -100/.9 = -111 < -100$
Interpretation: • Investr worse off with risk so pays less PV to buy asset. • Risk-averse investr doesn’t like positive risk.	Interpretation: • Investr better off with risk so pays more PV to buy asset. • Risk-averse investr likes negative risk.	Interpretation: • Investr better off with risk so accepts less PV pmt to take on risk. • Risk-averse investr likes positive risk.	Interpretation: • Investr worse off with risk so requires larger PV pmt to take on risk. • Risk-averse investr doesn’t like negative risk.
Examples: Invest “long” in stocks, bonds, or real assts.	Examples: Put options on stocks, bonds, or real assets. Insurance.	Examples: Construction costs. Borrow money. “Short” positions.	Examples: Invest “long” in “out-of-the-money” asset without full ltd.liability. (e.g., copper mine when copper price low & mine subject to fixed costs.)
Replication: Twin assets. Synthetics (index-linked).	Replication: Combi short stock long bond.	Replication: Short twin assets. Issue bond. Synthetics.	Replication: Combi long stock short bond

Summarizing the cases for one investor:

The above described four different types of risk situations are summarized in the table above. And the investor's risk-averse utility function that underlies the valuation of the cases is portrayed in the figure on the preceding page. The utility function describes the investor's preferences, her amount of satisfaction from different levels of wealth. The function has a curved, concave shape, meaning that although it is always upward sloping (more wealth is preferred to less wealth), it "bends" downward, its slope is ever declining (positive first derivative, negative second derivative, as a function of wealth). In combination with the proposition that under uncertainty surrounding the future the investor maximizes the probabilistically expected utility (that is, the mean utility over the possible future outcomes), the concave utility function results in classical "risk-averse" behavior, in which the investor must be "paid" in order to take on more risk. Holding expected outcome constant (at a point on the horizontal axis), the investor always prefers a mutually exclusive alternative that has less risk in its future outcomes (less range along the horizontal axis), and this is true whether the outcomes are positive or negative.

From one investor to the market as a whole: Extending the SML to "negative risk"...

Now suppose the investor we've been talking about in the above examples, the one on whose wealth we have been considering the impact of the outcomes of the subject investment in the four cases of different types of risk, is an investor who is "representative of the market". That is, this investor is typical of the "average" investor in the market, or more to the point, typical of the marginal investors in the capital markets, the investors who determine the market prices of all assets in general equilibrium. Then the discount rate that produces the described results in terms of how much the investor is willing to pay (or must be paid) for the subject asset, the discount rate that is consistent with the concave (risk-averse) utility function in the above graph, will be the market discount rate. In other words, it will be the market's required return (that produces the equilibrium market values of all the assets). In other words, it will be the relevant market opportunity cost of capital (the "OCC") for the asset or claim in question.



In such circumstances, we can directly translate this “representative investor’s” asset valuations and implied discount rates, as the OCC, to the “Security Market Line” (SML) depicted in the figure above. This is the same, classical SML you have probably seen before (e.g., in descriptions of the “CAPM” where the horizontal axis is measured by “beta”), only now the portrayed SML is extended into “negative risk” territory. But this extension really makes no difference. It is still the same SML, and it still means the same thing. It keeps the expected return risk premium per unit of risk constant across all assets on an ex ante basis. Only now, we recognize that risk can be “negative”, in the sense described above, of corresponding to a negative OCC.³ The SML is positively sloped because the (representative) investor (i.e., the market) is risk-averse (reflects concave utility).⁴ And the SML is a straight line in order to keep the risk premium per unit risk

³ Recall that we’ve ignored for simplicity the case where the risk premium is negative but the OCC is positive (only less than the riskfree rate). Such a case is not qualitatively different from Case 1 or Case 3 that we have already described (depending on whether the expected outcome is positive or negative), and the $0 < \text{OCC} < r_f$ case is included in the SML depicted in the figure above. However, this case does not fit the terminology we have been using. A negative expected return risk premium ($\text{OCC} - r_f < 0$) would generally be called “negative risk”, yet the positive OCC makes it behave as in our Case 1 or Case 3 descriptions. It is a rare circumstance in which risk is negative but offset in present valuation and future return expectation by the time value of money which is always positive.

⁴ Suppose the market were not risk-averse, as though the representative investor’s utility function were convex instead of concave. Then the certainty-equivalence discounting present value formula that we are fundamentally

constant. This is nothing more (or less) than the “Law of One Price”, as applied to the price of risk in terms of the investor’s (the market’s) expected return risk premium. The Law of One Price is a defining condition of equilibrium in markets, and from a normative perspective it can be viewed as a condition of “fair” pricing or valuation of assets, at least relative to one another. (The “Law” simply states that the same thing can’t trade at different prices at the same time in the same market.)

A methodological “bottom line”: Use certainty-equivalence discounting...

One “take-away” we can get from this discussion is that you can’t always use risk-adjusted discount rate discounting. Indeed, once we leave the familiar territory of just valuing “underlying assets” (stocks, bonds, real assets), the OCC becomes a rather slippery and tricky concept and one that is not necessarily very useful for valuing assets or claims. OCC varies according to the state of the world, often cannot be determined exogenously which means it must be backed out from an already-computed present value, and sometimes the OCC does not exist at all (even for a simple 1-period asset, whenever the $PV=0$, so this is not the same mathematical problem as besets the IRR). And other things being equal (namely, future expected outcome), a negative OCC can either be a “good thing” for the investor (if expected outcome is positive) or a “bad thing” (if expected outcome is negative).

On the other hand, you can always use certainty-equivalence discounting (section 16.3.2 of Chapter 16), whether it be to value plain vanilla underlying assets, or “exotic” claims and contracts and deals and operations. In the territories of “*financial engineering*” and of “*real options analysis*”, certainty equivalence discounting is the methodology to use for valuation. But the OCC (discount rate) still has intuitive appeal, as a way to measure and describe the amount of investment risk. You just have to “back it out” (when that is possible).⁵

using in our option evaluation methodology would still work. It would still be the correct formula and methodology. But the SML whose ex ante risk-premium per unit of risk that formula is reflecting would then be negatively sloped instead of positively sloped. (The underlying asset risk premium on the right-hand-side of the PV formula for the derivative would be negative.) This would reflect a weird world where more risky underlying assets command higher prices, unlike the real world.

⁵ And sometimes the OCC is useful to help determine the future expectation to which certainty equivalence discounting should be applied. For example, in pricing futures or forward contracts the “beta” and the SML regarding the underlying commodity or index can be useful to form the expectation about the future value of the underlying asset, and thereby help to set the initial “price” of the contract (which will be paid in the future).