



The Psychology of the Stock Market Professor Alex Edmans

22 September 2020

The Efficient Markets Hypothesis

The Efficient Market Hypothesis (EMH) is one of the most famous theories in finance, with Eugene Fama winning the Nobel Prize in Economics for his work on this topic. According to the theory, the price of any security reflects all information that's relevant for determining its value.¹

Let's take Apple stock as an example. What information is relevant for determining its value? Examples might be the state of the U.S. economy, customer reviews of the iPhone versus its competitors, and market consensus on the quality of Apple management. If the market is efficient, Apple's stock price will react quickly to any releases of relevant information. For example, if the U.S. economy announces good unemployment numbers, Apple's stock price should rise. If Apple's CEO has a health scare, its price should fall.

The implications are profound. If the EMH is true, then it means that investors can't make money trading on information. Any information that an investor might want to use will already be in the price. Indeed, the EMH doesn't seem to be just an abstract theory – many people believe it holds in practice. Historically, *mutual funds* (investment funds that hold multiple stocks) were typically *actively-managed*. There's a fund manager who chooses which stocks to buy and which to avoid, based on information. All this takes time and effort, and costs money. The fund needs to hire a manager and her team of analysts, and they need to gather information and crunch the numbers. So, in return the fund will charge its investors an *annual management fee* of, say, 1% of fund value per year. But, in recent years, there's been a rapid growth in *index funds*. Here, there's no actual manager – instead, a computer just buys all shares in an index. As a result, they're much cheaper, charging 0.1% per year or even less. This growth is underpinned by the belief that there's no point paying for a fund manager, as any information she'll be using is already in the price.

But not only is the EMH one of the most famous theories in finance – it's also one of the most controversial. Bob Shiller and Dick Thaler have themselves won the Nobel Prize in Economics for their work contradicting it. They are scholars of *behavioural economics*, a field which studies the impact of human psychology on decisions. Human investors may act differently from the predictions of the EMH, for two reasons.

1. They may not have relevant information. The crucial word in the definition of the EMH is "all". It means that every bit of information that's relevant in determining Apple's value – even if it's only tangentially relevant – should be reflected in Apple's stock price. For example, let's say that there's the possibility of a military coup in Venezuela. That's relevant information, since Apple might sell 0.5% of its iPhones to Venezuela – so if a coup occurs, this would hit the Venezuelan economy and reduce iPhone sales. But the main investors who trade Apple stock might be based in the U.S., and don't know what's happening in Venezuela.

¹ Eugene F. Fama (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25, 383-417.

2. They may not understand the information. Even if investors have all relevant information, they may not know how to process it. Some may overreact to it, others may underreact to it. For example, even if an investor knew that a military coup was likely, she may not be able to forecast how much it will hit iPhone sales and for how long.

The difference between these two reasons is like the difference between poker and chess. Poker is a game of imperfect information, where people make mistakes because they can't see their opponents' cards. Chess is a game of perfect information. But even though a player can see all the pieces, he may still make mistakes by not being able to process the information correctly. This distinction is important because it means that, even in a big data world where information is becoming increasingly available, markets may still be inefficient.

Defenders of efficient markets will acknowledge that people make mistakes. But they'd argue that mistakes shouldn't affect prices because they're random. If Andrea is overoptimistic about a stock and buys too much, and Bimal is pessimistic and sells too much, they'll cancel each other out. But advocates of behavioural finance argue that what causes mistakes is human psychology. Because Andrea and Bimal are both humans, they're affected by the same psychological biases. They'll make mistakes in the same direction and reinforce rather than neutralise each other.

Let's look at two major psychological biases that humans have, and see how they affect the stock market.

Overreaction

One common bias is *overreaction*. People overextrapolate from small patterns in the data. If you go to a sports stadium this weekend, and see a manager lose for the third straight game, fans will call for his head. More precisely, they're over-attributing the stream of losses to the manager's ability rather than luck. Even a high-ability manager, who wins 50% of his games, will still have streaks of three losses due to chance – just as a fair coin will still have streaks of three tails. As Nassim Taleb explained in a book of the same name, people are “fooled by randomness”.

Now translate this from the sports stadium to the stock market. A tech company announces an increase in sales. This increase is due to luck (good economic conditions) but investors mistakenly attribute it to product quality. They think sales will continue to rise and pour into the stock, causing its price to soar. The soaring stock price triggers its own overreaction as investors expect the price rise to continue.

If markets indeed overreact, then this implies a profitable trading strategy, known as *reversal*. One of Thaler's early papers, with Werner de Bondt, sorted stocks according to their performance over the past 36 months.² The stocks that had done the best, they called the “winners”, and the ones that had performed the worst were dubbed the “losers”. They found that, over the next 36 months, past “losers” outperform past “winners” – there's a reversal of fortunes. A recent “winner” may have enjoyed mildly good news, but investors mistakenly thought it was very good news and driven its price up too high. That's why it subsequently underperforms. So, in contrast to the EMH, there's a profitable trading strategy – buy past losers and sell past winners. This appears to imply a huge level of market inefficiency, as it seems that you can make money without even knowing the name of a company or what industry it's in. All you need to do is look at its past stock price performance to figure out whether to buy or sell it.

Underreaction

² DeBondt, Werner and Richard H. Thaler (1985): “Does the Stock Market Overreact?” *Journal of Finance* 40, 793-805.

A second common bias is *underreaction*. People are too slow to update their beliefs in the face of new data. This may be due to *confirmation bias* – the tendency to reject any data that contradicts their prior beliefs, and accept data that supports them. If you’re a climate change denier and you hear of a study that concludes that climate change is a hoax, you instantly accept it. If you see research claiming that climate change is real, you read it trying to poke holes in every argument you encounter. So your beliefs don’t change. (For more detail on confirmation bias, please see my TED talk, “What to Trust in a Post-Truth World.”)

Now translate this from science to the stock market. It’s 1981 and sales of the film company Kodak have just crossed \$10 billion per year. But Sony has just released the Mavica, the first ever electronic camera (a precursor to the digital camera). How did investors respond? Many didn’t respond at all. Kodak was one of the world’s largest companies – it seemed unfathomable that they’d be threatened by such an untried invention. So the stock price stayed high.

If markets indeed underreact, then this also implies a profitable trading strategy, known as *momentum*. Like De Bondt and Thaler, Jegadeesh and Titman also sorted stocks into winners and losers, but now on performance over the past 6 months, not 36 months.³ Instead of reversal, they found evidence of momentum: past winners continue to be winners, while recent losers keep dropping.

Combined with the findings of De Bondt and Thaler, this suggests that the market underreacts in the short-term (implying short-term momentum) but overreacts in the long-term (implying long-term reversal). If a company announces good news which causes its fundamental value to increase by 20%, its stock price will initially underreact. So if you buy a stock that’s risen over the past six months, it has further upside. But, after a while, the stock price overshoots. So, a stock that’s risen over the past 36 months has gone up too much and you should sell it.

Underreaction to Events

So there’s evidence that the market underreact in the short-term. We can drill down further by investigating *what* the market underreacts to.

A famous paper found that the market doesn’t even fully react to one of the most important and visible corporate events: quarterly earnings announcements.⁴ Every three months, a company announces its earnings – an event that the market pays substantial attention to. Now even if earnings rise relative to the last quarter, that might not be “news” if the increase were predictable. The only new information that should move the stock price is the unexpected portion of the new announcement. How do we strip out this component? Before the announcement, equity analysts (such as Goldman Sachs and Morgan Stanley) will forecast what they think quarterly earnings will be. The average forecast is known as analyst consensus. The new information in the earnings announcement is the actual earnings minus the analyst consensus, known as the *earnings surprise*. How does the market respond to earnings surprises?

If earnings end up above expectations – a *positive earnings surprise* – the stock price should jump on the announcement but then be flat afterwards. But this isn’t the case – it continues to drift up afterwards. This *post-earnings announcement drift* seems extremely puzzling. Since earnings announcements are such important events, it seems bizarre that the market doesn’t respond to them. If you told your friend to short a stock because, last month, it had a bad earnings announcement, she might think that’s a dumb idea since it’s old news – and very public old news. But the evidence shows that the idea may be sound, since the stock price continues to drop even after 30 days.

³ Jegadeesh, Narasimhan and Sheridan Titman (1993): “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency.” *Journal of Finance* 48, 65-91.

⁴ Bernard, Victor L. and Jacob K. Thomas (1989): “Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?” *Journal of Accounting Research* 27, 1-36.

How can the market underreact to even an event as visible as a quarterly earnings announcement? Let's go back to the two reasons for why investors may mistakes. The second was that they may not respond correctly to information. Even if investors saw the earnings announcement, they may underreact to it due to confirmation bias. The first was that they may not notice the information. Even though earnings announcements are indeed major events, professional investors may have to cover hundreds of stocks. A study shows that post-earnings announcement drift is even greater when a company announces earnings on days when many other companies are announcing earnings.⁵ The interpretation is that they're distracted by those other earnings announcements, and don't have capacity to process an additional one. Another paper finds that post-earnings announcement drift is greater when companies announce earnings on Fridays.⁶ The interpretation is that investors are wrapping things up for the weekend – if a company is announcing earnings on Friday afternoon, the investor isn't going to stay in the office and rerun her valuation model to take it into account.

Sentiment

So far, we've argued that the market might be inefficient by showing that it doesn't respond to information that it should – such as past prices or earnings announcements. A second way to do so is to show that it does respond to information that it should not.

Yet some studies have suggested that sentiment may affect the stock market. One shows that the market falls when the clock changes from Spring to Fall or Fall to Spring.⁷ Their interpretation is that clock changes mess up traders' sleep patterns and they go to work in a depressed mood. Another finds that the stock market declines on cloudy days – and conducts the analysis across several countries to check robustness.⁸

But you might still have concerns of spurious correlation. Are clock changes and weather strong enough to affect how investors trade? Certainly, sleep is important for your mood, but traders are pretty good at getting by on little sleep. After a late night, they typically have another shot of espresso. Similarly, weather is not correlated across a country. It could be sunny in New York, but cloudy in Chicago, and it might be Chicago fund managers who are driving stock prices even though the stock exchanges are located in New York.

In my PhD thesis at MIT, I thus (with coauthors Diego Garcia and Oyvind Norli) investigated the effect of international football results.⁹ Studies show that sports defeats lead to serious effects – people suffer heart attacks, commit suicide, and even murder each other – effects that are unlikely to be neutralized by even the freshest espresso. Moreover, international defeats affect the whole country in the same way – when England lose, the whole of England is upset, regardless of what city you're in. (That's why we couldn't study the Premier League – if Chelsea win and Manchester United lose, then some Brits are happy and others are unhappy, so it's hard to predict what will happen to the overall stock market).

⁵ Hirshleifer, David, Sonya Seongyeon Lim and Siew Hong Teoh, (2009): "Driven to Distraction: Extraneous Events and Underreaction to Earnings News" *Journal of Finance* 64, 2289-2325.

⁶ DellaVigna, Stefano and Joshua M. Pollet (2009): "Investor Inattention and Friday Earnings Announcements" *Journal of Finance* 64, 709-749.

⁷ Kamstra, Mark J., Lisa A. Kramer and Maurice D. Levi (2000): "Losing Sleep at the Market: The Daylight Saving Anomaly" *American Economic Review* 90, 1005–1011.

⁸ Hirshleifer, David, and Tyler Shumway (2003): "Good Day Sunshine: Stock Returns and the Weather" *Journal of Finance* 58, 1009–1032.

⁹ Edmans, Alex, Diego Garcia and Oyvind Norli, (2007): "Sports Sentiment and Stock Returns" *Journal of Finance* 62, 1967-1998.

We found that stock markets typically fall by 0.5% the day after the national team is eliminated from the World Cup. 0.5% may not seem a lot, but applied to the UK stock market, that's nearly £10 billion in a single day. As expected, the effect is stronger in the World Cup than the European Championships, Copa America or Asian Cup; stronger in elimination stages than group games and qualifiers; and strongest of all in England, France, Germany, Italy, Spain, Argentina and Brazil. The effect is astonishingly robust across the 39 countries. We also find a loss effect in international rugby, cricket and basketball – with an especially large effect for cricket in South Asian countries.

There's surprisingly little effect for wins, which could be for many reasons. One is that fans are systematically overconfident about their team's prospects – every other year for 54 years, England fans have thought that this would be “the year” we'd win a major championship, always to be disappointed. If fans go into each game expecting they'll win, there's little effect if they do win, but they become depressed if they lose. Another is the asymmetry of the competition format: winning an elimination game merely sends you into the next round, but losing leads to instant exit.

Astute readers may think there's a profit opportunity here: investors should sell stocks in both countries before a game, as the winning team will show no effect but the losing country's stocks will fall. However, international sports matches probably don't occur often enough to have a portfolio fully dedicated to trading them. The goal of the study was not so much to find a profitable trading strategy (that's the objective of hedge fund research, not academic research) but to show that markets are inefficient.

Relatedly, you might think – is this a serious academic study? Who cares about the effect of football on the stock market? But the goal of the research wasn't to investigate whether football affects the stock market, but a much broader question – is the stock market driven by fundamentals or sentiment? We study football as it's a variable that affects sentiment but has little effect on economic fundamentals such as profits and dividends.

© Professor Alex Edmans, 2020

Some of this summary is adapted from *Principles of Corporate Finance* by Richard Brealey, Stewart Myers, Franklin Allen, and Alex Edmans (14th edition, forthcoming in 2022).