



## The Mistakes Investors Make Professor Alex Edmans

8 DECEMBER 2020

### Introduction

Lecture 1, The Psychology of the Stock Market, showed that the stock market is driven by psychological biases, not fundamentals. Lecture 2, Hidden Investment Opportunities, showed that the market overreacts to salient, tangible information and underreacts to non-salient, intangible information. In this lecture, we'll drill down even further into the mistakes investors make. In particular, in addition to looking at the behaviour of the overall market, we'll go one level deeper and study the behaviour of individual investors.

### Overconfidence

One extensively-documented psychological bias is overconfidence. If you ask people whether they're above average in their driving skill, ability to get along with others, sense of humour – pretty much anything – 90% of them say yes. Since only 50% of people can be above average, this is a clear sign of overconfidence.

What does overconfidence mean for trading behaviour? It implies that household (i.e. amateur rather than professional) investors will trade too much. Often investors may have a bright idea about why they should buy or sell a stock, not realising that these insights might already be in the stock price. While the Efficient Markets Hypothesis (EMH) doesn't hold perfectly (see Lecture 1), it's an important benchmark. An investor should only trade if she has a good reason to believe that the EMH doesn't hold – why her insights might be missed by the thousands of professional investors who will be studying the stock on a full-time basis.

The challenge with testing this implication is that individual investors' trading behaviour is private information. You can observe total trading volume on the market, but that doesn't tell you who's buying or who's selling. By law, large U.S. institutional investors have to disclose their positions every quarter, but there's no such requirement for households. However, Professor Terry Odean obtained a dataset of individual investor trades from a large anonymous brokerage in the mid-1990s, which he then generously shared with other researchers, allowing for an extremely influential stream of research on individual investor behaviour.

Terry's work, mostly co-authored with Professor Brad Barber, found that individual investors lose significantly from trades. The stocks they buy underperform the stocks they sell by 0.23%/month, and that's even before taking transactions costs into account.<sup>1</sup> Using a similar dataset but from a later period, households in aggregate earned 18.4% per year, actually slightly higher than the market which earned 17.9%/year (although the difference wasn't statistically significant).<sup>2</sup> But after deducting

---

<sup>1</sup> Odean, Terrance (1999): "Do Investors Trade Too Much?" *American Economic Review* 89, 1279-1298.

<sup>2</sup> Barber, Brad and Terrance Odean (2000): "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors". *Journal of Finance* 55, 773-806.

trading costs, the return was 16.7%/year. Households that were frequent traders earned only 11.4%/year, suggesting that households lose money by trading. The best thing to do is just to stay put, rather than being overconfident in your own insights and trading on them.

Moreover, the above figures are only returns; they don't take risk into account. Terry and Brad also found that the average investor only holds 4 stocks, so is substantially underdiversified. Likely investors believe they have unique insights about these 4 stocks, hence putting all their eggs into four baskets, but the evidence suggests that they'd be better off diversifying.

Interestingly, the overconfidence hypothesis has cross-sectional predictions (i.e. between different investors). There's abundant psychological evidence that men suffer from overconfidence more than women. Indeed, men trade 45% more than women. An individual trade by a man isn't worse than an individual trade by a woman (both lose equal amounts of money), but because men trade more, they lose significantly more money each year.

Another cross-sectional test is between telephone and online traders. Psychological evidence suggests that overconfidence is exacerbated by the "illusion of control", the belief that you're in control of an event even when it's down to chance (e.g. people who throw the dice particularly hard when playing craps). Trading online gives you the illusion of control, since you can execute trades directly, rather than having to go through a telephone agent. Indeed, online traders trade more than telephone traders, and do worse. Moreover, there are time-series as well as cross-sectional predictions. Taking an individual trader and tracking her trades after moving from telephone to online trading, she trades more and loses more money.

### Attention

Under the EMH, investors should search over every possible stock – including those overseas – when deciding what to buy. But in the real world that's impossible, particularly for household investors – they have limited time, and so can only focus on a subset of stocks. This is an example of *bounded rationality* – people don't make perfect decisions because they can't take all relevant information into account.

Since investors can only have a few stocks on their radar screen, how do they decide which ones end up there? A plausible answer is stocks that have caught their attention – perhaps by (1) being in the news, (2) having a large price movement on the previous day, or (3) if there's a lot of activity around the stock (which manifests in lots of people trading the stock, i.e. a high trading volume). Note that attention is far less likely to be important for sales. When an investor decides what to sell, she can only sell the shares she currently owns, since most investors don't short-sell. Indeed, a study finds that individual investors are net buyers of attention-grabbing stocks (where "attention-grabbing" is defined using the three above ways).<sup>3</sup> These stocks subsequently underperform. However, this wasn't the case for professional investment firms, who can short-sell, can search over more stocks (as they have many employees), and may already have screens to limit their choice set (e.g. one fund might focus on healthcare stocks) - removing the need to limit it further by focusing on attention-grabbing stocks.

While that study focused on attention which comes from external events, attention can also arise internally from an investor's own portfolio. A separate paper uncovered the *rank effect* – that investors

---

<sup>3</sup> Barber, Brad and Terrance Odean (2008): "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors". *Review of Financial Studies* 21, 785-818. The finding that the bought stocks subsequently underperform is not in the published version of the paper but was in the working paper version.

are most likely to sell the biggest winner and the biggest loser from their portfolio than stocks in the middle; the sold stocks subsequently outperform.<sup>4</sup> When an investor logs onto her portfolio, she's most likely to notice the biggest winner and loser, and then decide whether or not to sell it.

Interestingly, not only do household investors exhibit the rank effect, but mutual funds too. Moreover, smart investors can exploit the fact that the prices of some stocks may have dropped excessively due to attention-based selling. You can look at a mutual fund's disclosure of its portfolio positions to figure out which stocks are best- and worst-ranked, and thus predict which stocks are likely to be sold. Buying such stocks 70 days later (by which time they are likely to have been sold) earns 1.36%/month for the worst-ranked stocks and 0.19%/month for the best-ranked. The much greater returns for worst-ranked stocks is because mutual funds sell them much more heavily than best-ranked stocks.

### Incorporating Irrelevant Information

Several individual investor mistakes can be classified around paying attention to information that's irrelevant. Here are three examples:

#### 1. The Original Purchase Price

In a world of no taxes, the *only* factor that should influence whether an investor sells a stock is the current price compared to her estimate of its true value. The original purchase price doesn't matter – it's a “sunk cost”. If you bought a stock at £100 and its price is now £90, but you think it's worth only £85, you should sell it. It doesn't matter that you're making a loss of £10; if you think it's heading further south to £85, not selling will cause your loss to widen to £15.

If there are taxes, the purchase price will matter, in that investors should be less willing to sell winners than losers. If you sell a winner, you're liable to capital gains tax, but if you sell a loser, you don't get a capital gains tax rebate (unless you have winners against which to offset these tax losses).<sup>5</sup>

However, the evidence is that investors do exactly the opposite – they sell winners and hold onto losers.<sup>6</sup> This behaviour, known as the *disposition effect*, is costly – not only because of taxes, but also because recent winners tend to outperform recent losers over the next six months. (This is the *momentum effect* discussed in Lecture 1, The Psychology of the Stock Market). So, the stocks that they sell end up outperforming the stocks that they retain. One potential explanation for such behaviour is *realisation utility* – investors get a jolt of utility, or happiness, when they realise a gain, and suffer a jolt of pain when they realise a loss.<sup>7</sup> It's hard to take a loss – which is why gamblers also “throw good money after bad” in a casino, trying to leave it at least breaking even.<sup>8</sup>

#### 2. Unrelated Past Events

<sup>4</sup> Hartzmark, Samuel M. (2015): “The Worst, the Best, Ignoring All the Rest: The Rank Effect and Trading Behavior.” *Review of Financial Studies* 28, 1024-1059. The finding that the sold stocks subsequently outperform is not in the published version of the paper but was in the working paper version.

<sup>5</sup> In the UK, there may be tax reasons to sell winners – investors have a capital gains tax allowance each year, and they may wish to sell winners to make use of this allowance. However, there is no such allowance in the US, where these studies were conducted due to data availability.

<sup>6</sup> Odean, Terrance (1998): “Are Investors Reluctant to Realize Their Losses?” *Journal of Finance* 53, 1775-1798

<sup>7</sup> Barberis, Nicholas and Wei Xiong (2012): “Realization Utility.” *Journal of Financial Economics* 104, 251-271.

<sup>8</sup> One may think that loss aversion can explain the disposition effect, but this is actually not the case. See Barberis, Nicholas and Wei Xiong (2009): “What Drives the Disposition Effect?” *Journal of Finance* 64, 751-784

Another type of irrelevant information is unrelated past events. *Contrast effects* have been documented in many non-financial settings. For example, people judge crimes to be less severe after hearing about egregious crimes; a great karaoke performance makes the rest of the crowd less willing to step up to the mic. Applied to financial markets, it suggests that investors will compare the earnings announcement of a company today to earnings announcements yesterday, even if they were of unrelated companies. The hypothesis is that, if companies announced strong earnings yesterday, a company announcing strong earnings yesterday will seem less impressive. Indeed, a study finds that the more positive yesterday's earnings announcements were, the lower the return to today's earnings announcements – the market just isn't that impressed.<sup>9</sup> Interestingly, these lower returns hold *regardless* of how good today's earnings announcement is. Very positive earnings announcements are interpreted a little less positively. Very negative earnings announcements are interpreted even more negatively.

Indeed, there's a trading strategy that exploits these incorrect reactions. If yesterday's earnings announcements were positive, sell firms scheduled to announce earnings today; if they were negative, buy firms scheduled to announce earnings today. Such a strategy earns an impressive 15%/year – even when focusing on large firms, where transactions costs are low.

### 3. Base Prices

Consider these two scenarios:

A: A blender costs £50. It would cost £5 at a store 5 miles away. Would you go to that store to save £45?

B: A TV costs £1,000. It would cost £955 at a store 5 miles away. Would you go to that store to save £45?

Most people would answer Yes to A and No to B. The thinking is that a £45 saving is huge when a blender costs £50, so it's worth it to travel 5 miles to save 95% of the price. But, a £45 saving is small compared to a £1,000 TV, so it's not worth the 5-mile trip.

But the “base” price is irrelevant. Regardless of whether you're buying a blender or a TV, you save £45 by travelling 5 miles. So, the only consideration that matters is whether the time and travel cost of the 5 mile trip is worth £45. Here, it's the *absolute* saving that matters, yet people consider *relative* savings since they incorrectly take the base price into account.

### Ignoring Relevant Information

In contrast, other mistakes involve ignoring information that's relevant. Again, we'll look at three examples:

#### 1. Base Prices

Did you read that correctly? Yes, that's not a typo. I just said base prices are irrelevant (yet paid attention to); now I'm saying they're relevant (yet ignored)?

---

<sup>9</sup> Hartzmark, Samuel M. and Kelly Shue (2018): “A Tough Act to Follow: Contrast Effects in Financial Markets”. *Journal of Finance* 73, 1567-1613.

That's indeed the case – because whether the base price is relevant or not depends on the setting. Rather than a consumption (spending) setting, let's consider an investment setting. Again, consider two scenarios:

A: The stock price falls from £50 to £5.

B: The stock price falls from £1,000 to £955.

Which stock has performed worst? Both have gone down by £45. But here, it's not the *absolute* change that matters, but the *relative* change. Why? With the blender and TV example, you're buying one item in each case. But in an investment setting, you're not comparing one share of A with one share of B. If you had £1,000 to invest, you'd have either bought 20 shares of A (so your £1,000 falls to £100) or 1 share of B (so your £1,000 falls to £955). Clearly, A is the worse investment.

But many investors focus on absolute changes and ignore the base price. For example, the Android, iPhone and E-Trade stock apps all show the change in price in absolute terms, with no relative changes side-by-side. What's the implication? Let's say investors think that certain events have a fixed absolute effect on the stock price – for example, a new CEO is worth £10 per share. Then, that will be worth 20% for stock A and only 1% for stock B. As a result, stock A will be more volatile.

Indeed, a study finds that low-priced stocks are more volatile.<sup>10</sup> Importantly, this holds even after controlling for the size of the firm (to take into account the fact that smaller stocks are bigger). Even more convincingly, volatility rises after a stock split. If a company used to have 1 million shares outstanding at £50/share, and splits its stock so that it now has 2 million shares outstanding at £25/share, nothing has changed in terms of the company's fundamentals, but its volatility jumps.

## 2. Last Years' Financials

Many companies' sales are seasonal. For example, Borders bookstore used to sell most books in the fourth quarter of each year, due to Christmas sales. Indeed, out of its 63 quarterly earnings announcements from 1995-2010, the 14 largest were in the fourth quarter. It seems obvious that the market should take such seasonalities into account – but it doesn't. Instead, it's surprised by the unsurprising – on average, the stock price rose 2.27% upon fourth-quarter earnings announcements (thinking they're good news when they're simply what was expected), and fell 3.4% for earnings announcements in other quarters.

A study found that this happens systematically, in large scale. Indeed, investors can exploit these mistakes. If you buy companies with high earnings seasonality (i.e. which announced good earnings in the same quarter over the last five years) and sold companies with low earnings seasonality, you'd earn about 4% per year. Surprisingly, even professional equity analysts don't fully understand the effect of seasonality. While they correct for 93% of the seasonal shift in earnings, they miss an important 7% of it, meaning they're surprised when seasonal earnings are announced.<sup>11</sup>

## 3. Last Year's Non-Financials

A company's annual report contains not only financial numbers, but also discussion, e.g. of the company's strategy and key risks. These discussions have ballooned over recent years. The average

<sup>10</sup> Shue, Kelly and Richard Townsend (2020): "Can the Market Multiply and Divide? Non-Proportional Thinking in Financial Markets." Working paper.

<sup>11</sup> Chang, Tom Y, Samuel M. Hartzmark, David H. Solomon and Eugene Soltes (2017): "Being Surprised by the Unsurprising: Earnings Seasonality and Stock Returns" *Review of Financial Studies* 30, 281-323

10-K (a regulatory filing in the US) contained 15,000 words in 1995, but over 60,000 words in 2017. As a result, investors may not bother to read them, or be able to digest what's important.

However, there *is* a way of figuring out what's important – compare this year's text to last year's. That doesn't come naturally to investors – for full-year financials, a company often presents this year's numbers side-by-side with last year's but doesn't do so with the narrative discussion. But comparison is easy to do with computers. A study finds that, if you buy firms that don't change their reports significantly from last year, and sell firms that do, you earn 7%/year.<sup>12</sup> If you focus on the changes in the Risk Factors section, the results are an impressive 22%/year.

Why? In theory, the changes could be in either direction, so why does it seem to be that changes are nearly always bad news? Because regulation forces companies to disclose negative information. Shareholders have successfully brought class-action lawsuits against companies that have failed to disclose bad news. Thus, if companies change their reports from last year, it's likely that they learned bad news and fear they have to disclose it else they'll be sued. In contrast, short-sellers haven't successfully sued companies for failing to disclose good news, so it's rare that changes in annual reports are positive. (Of course, it's in companies' interest to disclose good news, but they typically do so immediately with press releases rather than waiting until the end of the year to do so in their annual report).

© Professor Edmans, 2020

---

<sup>12</sup> Cohen, Lauren, Christopher Malloy and Quoc Nguyen (2020): "Lazy Prices" *Journal of Finance*, forthcoming.