

“A MODEL OF INVESTOR SENTIMENT” (Authors: N. Barberis, A. Shleifer, R. Vishny)

The point of this note is to explain the ideas in the above research paper without using any mathematics or technical jargon (the original paper contains some of both). The intended reader is someone who is interested in economics and finance but who is not an academic researcher. I welcome your comments on the ideas below, whether you agree with them or not; and also on the write-up itself -- for example, please let me know if it is confusing, so that I can rework it.¹

I'll start with a short summary, and will then give the longer version.

SHORT SUMMARY

Historical data show that stocks with certain characteristics – for example, stocks with low price-to-earnings ratios, or stocks with good performance over the past six months – tend to perform particularly well. We argue that these patterns arise because some investors exhibit psychological biases in their thinking – specifically, two biases known as representativeness and conservatism.

LONGER SUMMARY

A lot of current research in finance is about something called the “cross-section of average stock returns.” In simple terms, it is about why certain kinds of stocks have better performance, on average, than certain other kinds of stocks. This paper presents one of the first “behavioral” models of the cross-section. In other words, it is one of the first papers to argue that several facts about the cross-section may reflect *less than fully rational thinking* on the part of some investors. (Two other well-known behavioral models of the cross-section are those of Daniel, Hirshleifer, and Subrahmanyam (1998) and Hong and Stein (1999)).

Some facts

Why is there a lot of research on the cross-section of average stock returns? In part, because there are many puzzling things about it. Here, in particular, are four puzzling facts that researchers are still trying to understand.

“Long-term reversals”

Suppose that you take historical data on U.S. stock prices and do the following exercise. Every three years, you form a group of stocks that had very *poor* performance over the previous three years (group A); and also, a group of stocks that had very *good* performance over the previous three years (group B). You then track the performance of the two groups over the *next* three years. What you will find, if you repeat this exercise

¹ This is a preliminary draft. Please do not quote or cite.

every three years using many decades of U.S. data, is that the stocks in group A do *better*, on average, than the stocks in group B. In other words, stocks with poor prior performance subsequently do *well*, while stocks with good prior performance subsequently do *poorly*. Why is this?

“The value premium”

Suppose that, every year, you form a group of stocks that have *low* P/E ratios, i.e. low ratios of price to earnings (or low ratios of price to cash flow, or of price to book value). Let’s call this group A. At the same time, you form a group of stocks that have *high* P/E ratios. Let’s call this group B. You then track the performance of the two groups over the *next* year. What you will find, if you repeat this exercise every year using many decades of U.S. data, is that the stocks in group A do better, on average, than the stocks in group B. Stocks with low (high) prices relative to some measure of the underlying firm’s fundamentals have good (poor) subsequent performance. Why?

“Momentum”

Suppose that, every six months, you form a group of stocks that had very *good* performance over the previous six months (group A); and also, a group of stocks that had very *poor* performance over the previous six months (group B). You then track the performance of the two groups over the next six months. What you will find, if you repeat this exercise every six months using many decades of U.S. data, is that the stocks in group A do better, on average, than the stocks in group B. In other words, stocks with good performance in the past few months subsequently continue to do well, while stocks with poor performance in the past few months subsequently continue to do poorly. Why is this?

(Some readers may wonder whether this “momentum” finding contradicts the “long-term reversals” finding. “Momentum” says that past winner stocks subsequently do better than past loser stocks. “Long-term reversals” seems to say the opposite. However, there is a crucial difference between the two results. In the case of momentum, we are following stocks that had good performance over the previous *six months*. In the case of long-term reversals, however, we are following stocks that had good performance over the previous *three years*. This difference in the time interval we use to measure past performance matters a lot, and is responsible for the contrasting results).

“Post-earnings announcement drift”

Suppose that, every quarter, you rank stocks based on the size of the surprise in their most recent earnings announcement (for example, based on how much the firm’s earnings beat, or failed to beat, analyst expectations). You form one group of stocks that had surprisingly *good* earnings news (group A); and another group of stocks that had surprisingly *bad* earnings news (group B). You then track the performance of the two groups over the next 60 days. If you repeat this exercise quarter after quarter, using many years of U.S. data, you will find that the stocks in group A do *better*, on average, than the

stocks in group B. In other words, stocks that just announced surprisingly good earnings *subsequently* do better than stocks that just announced surprisingly bad earnings. Why is this?

(To be clear: the stocks in group A obviously jump up in price when the good news about their earnings is released. But that's not the interesting part. The interesting part is that these stocks *continue* to perform well in the weeks that follow. Similarly, stocks in group B obviously drop in value when the bad news about their earnings is released. But what's more interesting is that these stocks *continue* to do poorly in the weeks after the announcement.)

So these are four important puzzles about the cross-section of average returns. Why do I call them "puzzles"? In general, if you want to explain why one set of stocks (e.g. stocks in group A) does better, on average, than another set of stocks (e.g. stocks in group B), *and* you believe that all investors are fully rational, then you only have one option: you have to argue that the stocks in group A are *riskier* than those in group B, i.e., that stocks in group A have higher average returns in order to compensate investors for their higher risk. The problem – and hence the puzzle -- is that, in none of the four cases listed above is there any clear evidence that the group A stocks -- the ones with the higher average returns -- are any riskier than the group B stocks.

A new approach

In this paper, Andrei Shleifer, Robert Vishny, and I try to understand the four puzzles described above using a different approach. Specifically, we argue that all four facts reflect less than fully rational behavior on the part of some investors – in particular, that they are the result of two prominent psychological biases: *representativeness* and *conservatism*.

Representativeness is a famous concept due to the psychologists Daniel Kahneman and Amos Tversky. One of its implications is that, after they see a sample of data, people can be too quick to draw inferences from the data. For example, if people see that a firm has posted high earnings growth for several quarters in a row, they may be too quick to decide that this is a firm whose *long-term* earnings growth is high (their mistake is to forget that even a firm with *modest* long-term earnings growth can post a few quarters of impressive earnings).

Shleifer, Vishny, and I point out that representativeness can generate the "long-term reversals" and "value premium" patterns I described earlier. Suppose that a firm posts a few quarters of surprisingly good earnings growth. As I noted above, if they suffer from representativeness, investors will be too quick to believe that the firm's *long-term* earnings growth is high. They will therefore bid the firm's stock price up too much, giving the stock a high P/E ratio. And given their overly high expectations, investors will, on average, be disappointed by the firm's subsequent earnings growth. On average, then, the stock's subsequent performance will be poor. In summary, we have an explanation for why a stock with a high P/E ratio generally earns a low return later on; in other words,

we have an explanation for the value premium puzzle. (A very similar intuition shows how representativeness can explain the long-term reversals puzzle).

Representativeness says that people are sometimes too quick to draw inferences from data. But psychologists have also found that, in some circumstances, people can be too *slow* to draw inferences from data – in other words, that they sometimes pay too little attention to the data and stick too much to their prior views. This evidence is sometimes given the label “conservatism”.

Shleifer, Vishny, and I note that conservatism can generate the “momentum” and “post-earnings announcement drift” patterns in the data. Suppose that there is a firm that you’ve thought about a lot – in particular, you’ve decided that its future earnings growth is going to be moderate – nothing too special. The firm then announces surprisingly good earnings. If you suffer from conservatism, you will react by saying: “I’ve already thought about this firm – its prospects are nothing special; I’m not going to change my mind based on *one* piece of information!” As a result, you only push the stock price up a little on the day of the announcement. And only in the next few weeks and months, once you realize your error, do you grudgingly push the stock price up to the right level. As a result, while the stock jumps up in price on the day of the announcement, it also continues to drift upward in the weeks after the announcement. This is exactly post-earnings announcement drift. (The intuition for how conservatism leads to momentum is very similar).

(Some readers will have noted that representativeness and conservatism are opposite effects, in some ways. In one case, people put too *much* weight on a data sample; in the other, they put too *little* weight on the data sample. It is therefore important to understand when one effect or the other dominates. Psychologists have not yet resolved this issue – it remains an open question. In our paper, Shleifer, Vishny, and I propose our own reconciliation. The idea is that if the investor sees a good earnings announcement *in isolation* – i.e. one that is not part of a *sequence* of good earnings – then conservatism dominates and the investor reacts too little. But if he sees a sequence of good earnings announcements, then representativeness takes over and he reacts too much).

References

Daniel, K., Hirshleifer, D., and A. Subrahmanyam (1998), “Investor Psychology and Security Market Under- and Overreactions,” *Journal of Finance* 53, 1839-1885.

Hong, H. and J. Stein (1999), “A Unified Theory of Underreaction, Momentum, and Overreaction in Asset Markets,” *Journal of Finance* 54, 2143-2184.