Course notes on behavioural finance

Jason Collins

2 August 2022

Table of contents

Overview

In this online book, I provide an introduction to behavioural finance, the study of how human decisions can drive market anomalies that would not be predicted by standard economic approaches to markets.

The book is based on part of a subject I teach in as part of UTS's Graduate Certificate and Masters of Behavioural Economics. The subject is called Behavioural Economics and Corporate Decision Making, and this material comprises two weeks of the subject. I have separated it from the rest of the material (provided in another book) as it is better framed as a separate body of knowledge. This material is taught through a mix on online-self guided learning, online seminars and in-person weekend workshops.

The Graduate Certificate and Masters is for post-graduates with no assumed prior knowledge of economics or behavioural economics. The subject on which this book is based is taken after introductory economics and behavioural economics units.

This book covers the following areas:

- **Financial markets foundations**: What are financial assets? What is the efficient markets hypothesis? How do we test the efficient markets hypothesis?
- **Behavioural finance**: What are some observed challenges to the efficient markets hypothesis? What human behaviours could be driving these phenomena?

Part I

Financial markets foundations

> "One day in December 2003, when Saddam Hussein was captured, Bloomberg News flashed the following headline at 13:01: U.S. TREASURIES RISE; HUSSEIN CAPTURE MAY NOT CURB TERRORISM.
>
> Whenever there is a market move, the news media feel obligated to give the "reason." Half an hour later, they had to issue a new headline. As these U.S. Treasury bonds fell in price (they fluctuate all day long, so there was nothing special about that), Bloomberg News had a new reason for the fall: Saddam's capture (the same Saddam). At 13:31 they issued the next bulletin: U.S. TREASURIES FALL; HUSSEIN CAPTURE BOOSTS ALLURE OF RISKY ASSETS.
>
> So it was the same capture (the cause) explaining one event and its exact opposite."
>
> Nassim Nicholas Taleb, The Black Swan: The Impact of the Highly Improbable

A feature of most daily news bulletins is a report on the performance of the local financial markets. "The ASX 200 is up 32 points." "Shares in Acme Corporation plunged on the release of poor retail sales figures in February." Unfortunately, as Nassim Taleb's story above suggests, these news reports are little more than exercises in storytelling, as journalists desperately attempt to attribute causes to what might just be noise.

But these financial markets become more interesting and worthy of examination when we start to think of them as the aggregations of thousands or millions of individual decisions.

In the book Consumer Financial Decision Making I discuss how financial decisions are made by individuals and households. Individuals use heuristics and rules of thumb in their decision making that, while often powerful tools, can also lead to systematic error.

In this book, we will examine how those individual decisions manifest in financial markets, and what is the aggregate effect of the various market participants.

I will start with the basic economic approach. I will then layer onto that a series of anomalies that the traditional approach has difficulty explaining, and follow with the research in behavioural finance that gives us a richer picture of markets.

1 Financial assets and markets

1.1 Financial assets

An asset is any resource owned by a person or business. Assets can be *tangible*, such as buildings and equipment, or *intangible*, such as goodwill, copyrights or patents. Tangible assets are anything that can be touched. Intangible assets lack that physical substance.

It is often hard to value an intangible asset. What is the value of the name "McDonalds"? What is the value of the copyright for my contribution on drums in a sound recording?

Financial assets are assets that get their value from a contractual right or claim of ownership. Financial assets include:

- Cash
- Bank deposits
- Stocks: A share of ownership in a company, entitling the owner to a share of the company's profits
- Bonds: A financial asset by which the bond issuer owes the bond holder a debt for which they are obliged to pay interest, the principal at a later date, or both
- Derivatives: A derivative is a contract that derives it performance from the performance of an underlying entity. Examples of derivatives include *futures*, which are an agreement to buy something at a predetermined price at a specified time in the future, or *options*, which give the option holder the right (but not obligation) to buy or sell an asset a specified price before or on a specified date.

Financial assets sit somewhere between tangible and intangible assets. In economics and law, financial assets are usually classed as intangible assets as they lack physical form. However, under accounting standards they are often classed as tangible assets as they are easily valued and should be included in company accounts. Their precise treatment varies across countries and legal systems.

1.2 Financial markets

A financial market is a market in which people trade financial securities (a trade-able financial asset). Securities include stocks, bonds, derivatives and commodities such as metals and agricultural outputs.

The type of financial market that you hear most about in the media are *exchanges*, which are organisations that facilitate the trade in financial securities.

In Australia, our best known exchange is the ASX. The ASX provides a facility where traders can buy and sell securities, such as stocks and derivatives. Buyers and sellers are matched, with prices on most of their markets publicly available in near real-time.

The beauty of public exchanges is that they provide a rich and detailed history of prices of stocks and other securities that we can examine. Heavily traded stocks might trade many times every minute, providing a highly detailed trajectory of their prices over time.

1.3 Valuing a financial asset

The value of a financial asset traded on a financial market has one easily accessible measure of its value: its price.

This does, however, raise the question of why a purchaser would want to pay that price for an asset. They must value the asset at the price or higher, or they would not purchase.

The value of a financial asset will ultimately be realised in the flow of payments that asset entitles the owner to in the future. For example, ownership of a stock will yield a stream of dividends to the owner until such time as they sell or the company is wound up.

As a result, one simple model of the value of a security is that the value equals the rationally expected present value of future payments. This is often called the "fundamental value" of a security.

Calculating this fundamental value requires a "discount rate", which reflects the stock's risk. Payments in the future are discounted by this discount rate. A challenge with any model is that you need to know what the appropriate discount for risk is, but this is not something we know.

2 The efficient markets hypothesis

The efficient markets hypothesis states that, if markets are efficient, then asset prices should reflect all available information. As such, prices should reflect the fundamental value of the asset.

When Eugene Fama (1970) first labelled the efficient markets hypothesis, he proposed three forms of market efficiency:

- **The weak form**: Prices reflect all information contained in historical returns. That is, beyond today's price, knowing the price history of a security should be of no value in predicting the future price of that security.
- **The semi-strong form**: Prices reflect all publicly available information. This includes past earnings, earnings forecasts, financial statements, stories in the media, economic conditions and anything else in the public domain that may be relevant.
- **The strong form**: Prices reflect all information, including that which is not publicly available. This includes private information held by insiders.

Practically, there is a cost to acquiring information, so the assumption that prices will always reflect *all* information is not reasonable. Therefore, a better working definition for the efficient market hypothesis is that prices reflect all information except that for which the marginal cost of acquiring the information exceeds the marginal benefit of acting on the information.

This last statement implies that, although there may be some information not yet incorporated into the price, no investor can generate excess returns. For example, when Malkiel (2005) examined mutual fund performance versus the S&P 500 index, around three quarters of funds underperformed on an annual basis and over 80% underperformed over 10-year periods. Similarly, in an analysis of household trading from 1991 to 1996, Barber and Odean (2000) found that the average household earns an annual return of 16.4 percent compared to a market return of 17.9 percent.

The efficient markets hypothesis is largely examined and tested using data from exchanges. The high levels of liquidity and the transparency of prices and dividends allow minute examination of whether the price is efficient. As we will see, however, determining whether a market is efficient is not an easy task.

3 Testing the efficient markets hypothesis

The most basic test of market efficiency is to examine whether the price of the asset equals its fundamental value. For example, does the value of a stock in a company equal the present value of expected future dividends?

3.1 The joint hypothesis problem

But as noted earlier, determining the fundamental value, or indeed any measure of the value of an asset, requires the flow of future payments to be discounted for risk. But what is the correct risk adjustment? As this is usually not known, this requires an assumption of the right model of risk.

Therefore, tests of the efficient markets have a challenge. If a test rejects the proposition that a market is efficient, is this because the efficient markets hypothesis does not hold in that market, or is it because we have not used the right model of risk? This question is known as the *joint hypothesis problem*.

Much of the debate about market efficiency that we are about to cover ultimately hinges on whether the assessments of market efficiency have used an appropriate model of risk.

3.2 The event study methodology

One classic test of the efficient market hypothesis is the event study methodology devised by Ball and Brown (1968). To test efficiency, look at a class of events that affect the price of stocks (e.g. earnings announcements) and examine the returns in the days leading up to the event, on the day of the event, and in the days after the event.

If markets are efficient, we should see a reaction to the event (e.g. good or bad news) in the days leading up to and including the event. There should be no further reaction on days after the announcement as an efficient market should react completely in response to the event as soon as the information is available.

The first papers to use event studies to examine market efficiency typically found evidence consistent with the weak and semi-strong forms of the efficient markets hypothesis. In one famous paper pioneering the event study methodology, Fama et al. (1969) examined stock splits, where a company is split into two. They found that prices immediately reflected all available information on announcement of the split, and no abnormal returns were achieved after the announcement (see the image below).

FIGURE 2a

AVERAGE RESIDUALS—ALL SPLITS

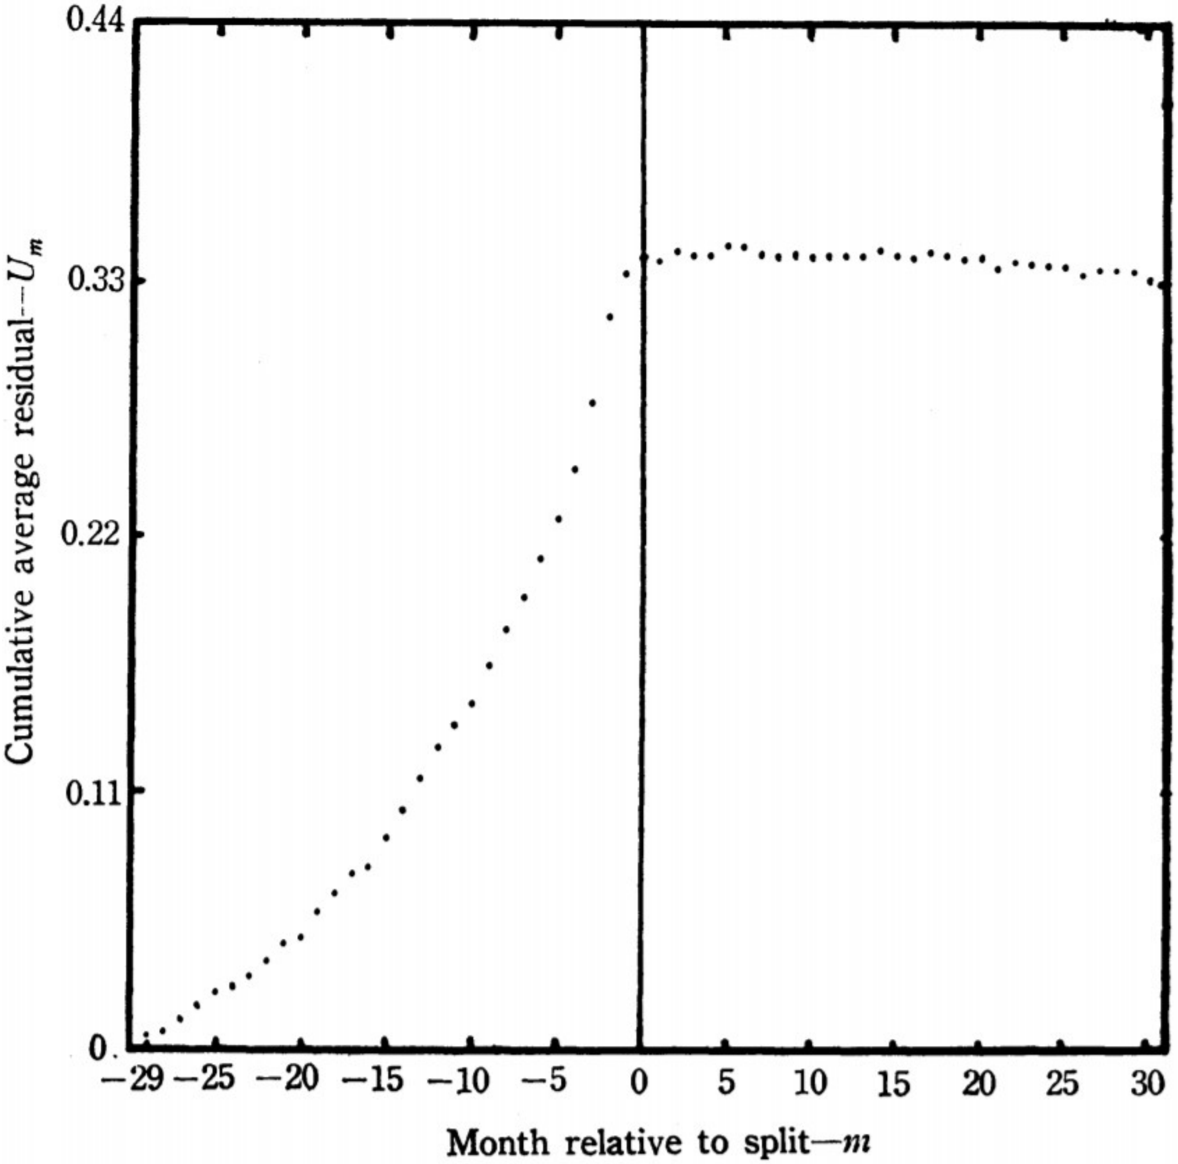

FIGURE 2b

CUMULATIVE AVERAGE RESIDUALS—ALL SPLITS

Part II

Behavioural finance

While the early days of the efficient markets hypothesis involved a parade of confirmatory studies, the last 30 years has seen a steady flow of theory and evidence that brings all three forms of the hypothesis into question. Recent studies have found many *anomalies*, empirical studies that bring market efficiency into question absent a better explanation.

Due to the joint hypothesis problem, “absent a better explanation” is a important qualifier. The tests of market efficiency are effectively a joint test of two hypotheses: market hypotheses and the particular model of risk used in the analysis. This means that any “anomaly” could be either evidence of market inefficiency or an inappropriate risk-adjustment technique.

The result is that every discovery of an anomaly is followed by a debate as to whether it is truly an anomaly and whether there is a rational explanation for it. That said, some anomalies have been hard to explain away.

In the remainder of this book I examine a series of anomalies. Three relate to anomalies concerning individual stocks: lagged reactions to events and announcements, momentum, and value stock out-performance. Three concern aggregate stock market performance: the equity premium puzzle, bubbles and excessive volatility. Finally, I will close with an examination of trader performance.

4 Lagged reactions to announcements

As noted earlier, the first papers to use event studies to examine market efficiency typically found evidence consistent with the market being efficient. However, since those early studies, improved methodologies and data have resulted in a steady increase in the evidence against market efficiency.

One example of a study questioning market efficiency was by Rendleman et al. (1982), who examined data around quarterly earnings announcements for around 1000 firms between 1972 and 1980. For each announcement they determined the extent to which the actual earnings varied from the forecast earning per share.

Unsurprisingly, they found that the market responds positively to positive surprises and negatively to negative surprises. That is consistent with the efficient markets hypothesis.

However, they also found continual drift in prices following the announcement, particularly announcements that involved a large surprise on either the upside or downside. This is inconsistent with market efficiency. It suggest either an under-reaction to the announcement, with the full reaction delayed, or an over-reaction in the long-term.

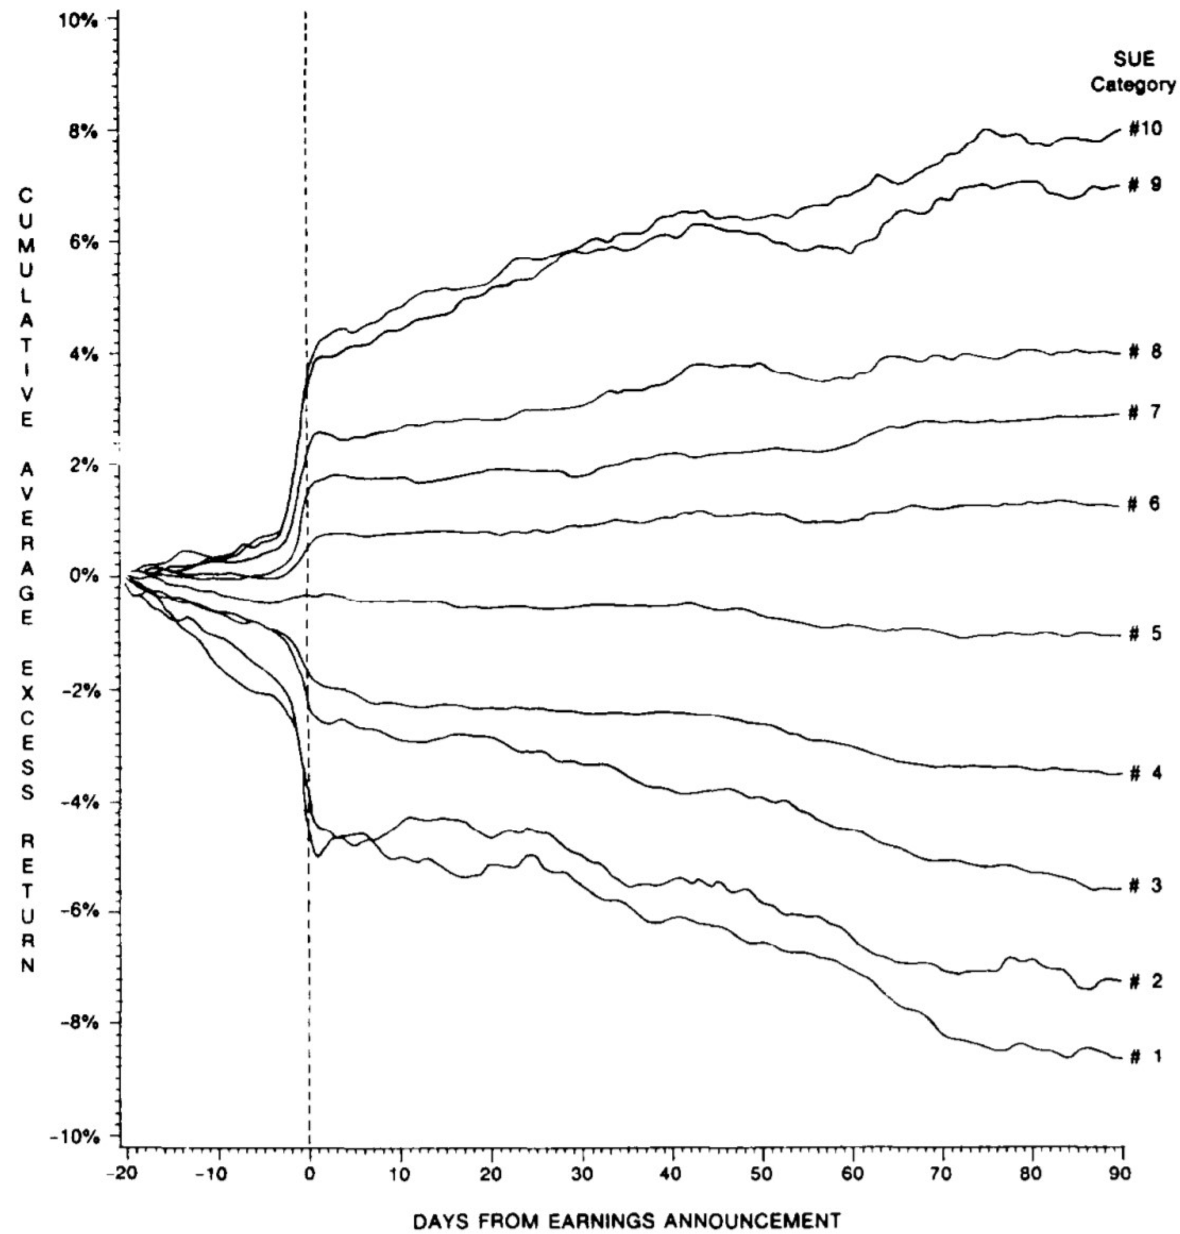

4.1 Potential explanations

In a review of the literature on post earnings announcement drift, Fink (2021) notes the following possible explanations:

- A stock's 52-week high provides an anchor. If a stock is close to the 52-week high, investors underreact to positive surprises. This leads to a subsequent upward drift. Conversely, investors under-react to negative surprises where stocks are far from their 52-week highs. This leads to a subsequent downward drift.

- The disposition effect is the tendency for investors to sell stocks that are in the gain domain relative to the purchase price and to hold stocks that are in the loss domain (Shefrin and Statman (1985)). The disposition effect will impede price adjustment following positive or negative news as investors will tend to sell after positive announcements and hold after negative announcements.
- Distractions or a complex environment can make it harder for investors to process earnings surprises. Post-announcement drift tends to be larger where there are more announcements on the same day or there is high market volatility.

5 Value stock out-performance

Value investing is a investment approach that involves buying stocks that appear under-priced based on some form of fundamental analysis. This is the bedrock of Warren Buffet's approach to investing.

Typically, that fundamental measure is based on an accounting measure such as earnings, cash flow or book value, rather than future growth. Value stocks are those stocks which have low prices relative to those accounting measures. Growth stocks are stocks with prices high relative to those measures as people anticipate high future growth.

There are many examinations of the performance of value stocks versus growth stocks. Fama and French (1992) grouped all stocks traded on the New York Stock Exchange, AMEX and NASDAQ into deciles based on their book-to-market ratio. Those in the highest book-to-market ratio decile, the value stocks, had an average return 1.5% per month higher than those in the lowest decile, the growth stocks.

5.1 A behavioural explanation

One possible reason for value stock out-performance suggested by Lakonishok et al. (1994) is that investors are making an error in the expectations of future growth as their expectations are excessively tied to past growth. They believe stocks with high growth in price now will have high growth in the future, despite the fact that future growth rates are highly mean-reverting (that is, stocks with high growth rates today tend to have more typical growth rates in the future).

Conversely, Fama and French (1992) simply argued that value stocks are riskier. (Recall the joint hypothesis problem.)

6 Momentum and reversal

The weak form of the efficient markets hypothesis states that returns should not be predictable using past returns. Despite this, there is evidence of medium-term momentum and long-term reversals.

6.1 Momentum

Momentum occurs when returns are positively correlated with past returns. There is reliable evidence that momentum exists over medium-term intervals of three to 12 months.

Jegadeesh and Titman (1993) grouped all New York Stock Exchange stocks traded between 1963 and 1989 into deciles based on their return over the previous six-months. They then calculated the average returns for each decile over the next six months. They found that the top decile of prior winners outperformed the bottom decile by an average of 10% (on an annual basis).

6.2 Reversal

De Bondt and Thaler (1985) ranked all stocks traded on the NYSE for every three years between 1926 and 1982 by their return over those three years. From these they formed a portfolios of 35 “winner” and 35 “loser” stocks and measured their performance over the following three years. They found that over the sample period the average return of the loser portfolio was 8% per year higher than the return of the winner portfolio.

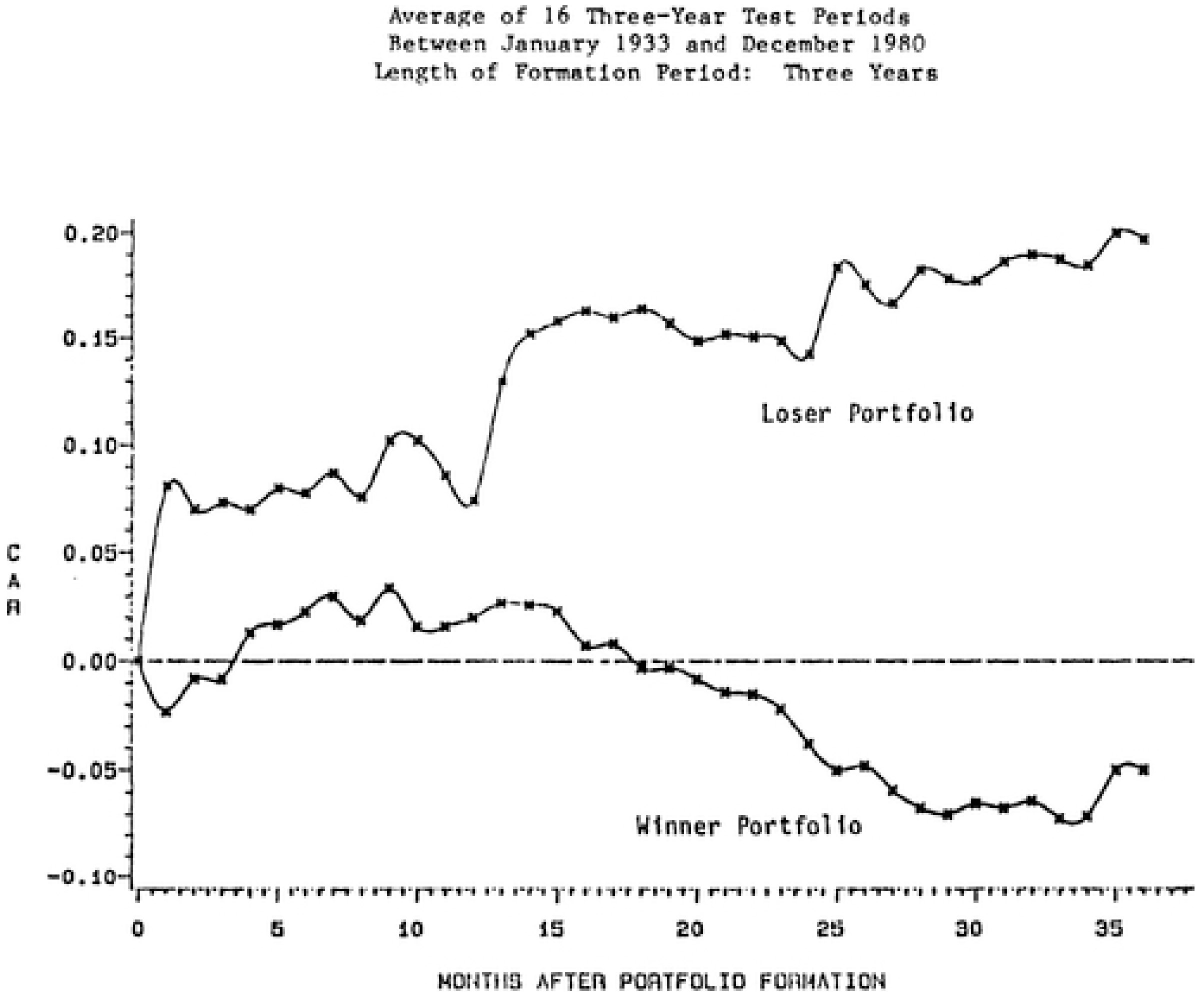

6.3 A behavioural explanation

The combination of momentum and then reversal suggests a combination of under-reaction then over-reaction, phenomena we have already seen in the explanations for lagged reactions to announcements and value-stock out-performance.

There are many attempts to explain this pattern. One proposed by Barberis et al. (1998) involves a combination of anchoring, by which people underweight new information relative to priors, and representativeness.

They proposed that when a company announces surprisingly good earnings, investors react insufficiently, meaning that the price does not increase as much as it should. This lower price, however, means that returns in subsequent periods appear good, creating post-announcement drift. This drift eventually causes people to overreact as they see the short period of good results as representative of the quality of the stock. This overreaction pushes the price too high. That high price then means that returns in following periods appear poor, leading to the eventual reversal.

7 The equity premium puzzle

The equity premium is the gap between the expected return on the stock market versus a portfolio of fixed-income securities (e.g. bonds). Since 1926 the annual real return for stocks in the United States has been about 7%, while the real return on Treasury bills has been less than 1%. (A Treasury bill is a short-period bond, issued by the United States Treasury, that pays a face value at maturity. They are bought at a discount to the face value to create a positive yield.)

This difference in yield might appear to be justifiable by the greater riskiness of stocks. If stocks are riskier, they should earn higher returns.

However, Mehra and Prescott (1985) examined this premium and argued that the size of the premium would require an implausible level of risk aversion. A reasonable level of risk aversion would result in an equity premium of around 0.1%. The size of the observed premium is hence a puzzle.

Two explanations for the equity premium puzzle are ambiguity aversion and loss aversion.

7.1 Ambiguity aversion

Ambiguity aversion is a preference for known risks over unknown risks. People don't just not know what the return will be. They also don't know what the potential distribution of returns is. They therefore require a greater premium for stocks than would be expected from risk aversion alone.

7.2 Loss aversion

7.2.1 Samuelson's bet

Consider whether you would accept either of the following bets.

- A 50% chance to win $200 and a 50% chance to lose $100?
- A sequence of 100 bets with a 50% chance to win $200 and a 50% chance to lose $100?

These bets relate to a famous exchange between Paul Samuelson and some lunch colleagues. Samuelson offered them a \$200 to \$100 bet that the side of a coin they specified would not appear at the first toss. One "distinguished scholar" responded:

> I won't bet because I would feel the \$100 loss more than the \$200 gain. But I'll take you on if you promise to let me make 100 such bets.

Samuelson (1963) showed that if a person would reject the first bet at any level of wealth, this pair of choices was not consistent with expected utility theory. The logic of his argument was as follows:

- The sequence of 100 bets could be thought of a sequence of 99 bets, plus a decision as to whether to accept one further bet.
- Given that this person will reject the single bet at any level of wealth, they will not accept this 100th bet.
- That leaves them with a sequence of 99 bets, which could also be thought of as a sequence of 98 bets, plus a decision as to whether to accept one further bet.
- Again, the additional bet is rejected.
- This logic is repeated until all bets are rejected.

Accordingly, Samuelson suggested that his colleague, if truly an expected utility maximiser, was making a mistake in accepting the bet of 100 flips given he refused the single flip. This is despite the fact that the sequence of 100 flips has an expected return of \$5 000, with less than a 0.05% chance of losing any money, and less than a 0.002% chance of losing more than \$1000.

Samuelson's conclusion would not change if someone had to accept the 100 bets as a single bet with that range of possible outcomes. Although the probability of loss is small, the variance of possible outcomes increases with the number of bets. This means the bet remains unattractive for a risk averse expected utility maximiser. And someone who would reject a win \$200, loss \$100 bet at any level of wealth would require absurd amounts of risk aversion, so much that this same person would also reject a 50:50 bet to win \$20,000, lose \$200 (Rabin and Thaler (2001)).

However, the response of Samuelson's colleague points to an alternative explanation for this pair of choices, that of loss aversion.

What if a person is loss averse? Suppose they have the following value function:

$$v(x) = \begin{cases} x & x \geq 0 \\ 2.5x & x < 0 \end{cases}$$

where x is a change in wealth relative to the status quo.

This loss averse person will turn down a 50:50 bet to win \$200, lose \$100:

$$
\begin{aligned}
V(x) &= 0.5v(\$200) + 0.5v(-\$100) \\
&= 0.5 * 200 - 0.5 * 2.5 * 100 \\
&= -25
\end{aligned}
$$

However, they would accept a sequence of two such bets, which has a distribution of outcomes of a 25% chance of winning \$400, a 50% chance of winning \$100 and a 25% chance of losing \$200.

$$
\begin{aligned}
V(x) &= 0.25v(\$400) + 0.5v(\$100) + 0.25v(-\$200) \\
&= 0.25 * 400 + 0.5 * 100 - 0.25 * 2.5 * 200 \\
&= 25
\end{aligned}
$$

Any longer sequence of bets has even higher positive value

Note, however, that this positive value for a sequence of bets only occurs if they do not have to watch the sequence of bets being played out. If they had to watch each consecutive flip, they would reject the bet as every individual flip has negative expected utility.

7.2.2 Explaining the equity premium puzzle

This story captures the intuition behind Benartzi and Thaler's (1995) explanation of the equity premium puzzle.

Suppose an investor has a choice between risky stocks, with an expected annual return of 7% and standard deviation of 20%, and a sure return of 1%. Like Samuelson's bet, the attractiveness of stocks to a loss averse investor will depend on both the time horizon of the investor and the frequency with which they evaluate the returns. If they monitor their portfolio frequently, they will often observe losses from stocks, which they feel with greater force than gains.

Suppose that one loss averse investor examines their portfolio every day. Since on a daily basis stocks go down almost as often as they go up, this investor will experience a lot of pain, making the stocks unattractive. Another loss averse investor only checks in on their portfolio once a decade. At that horizon, stocks have only a small probability of losing money, so will be much more attractive to someone who is loss averse.

It is a combination of loss aversion and a short evaluation period that will drive an investor to require a large premium for holding the risky option. Benartzi and Thaler call this *myopic loss aversion*.

8 Bubbles

On 10 March 1999, the technology firm-rich NASDAQ composite index closed at 2,406. On 10 March 2000, it had more than doubled to 5,048. By 9 October 2002 however, the index had crashed to 1,114.

It is commonly argued that the pricing of these firms comprising the NASDAQ was irrational or a "bubble". Bubbles are said to exist when prices are far above an intrinsic value based on fundamentals.

There is a large academic debate about what comprises a bubble. Do bubbles exist? Can a bubble be identified before a crash?

However, lab experiments provide a compelling case that, when many of the possible rational explanations for very high prices are eliminated through the experiment design, that bubbles do exist.

8.1 An experimental bubble

The first paper to explore bubbles in an experimental asset market was by Smith et al. (1988). This has now led to a large literature on the topic.

A typical bubble experiment involves an asset that is to be traded over a fixed number of periods. The asset pays a dividend on a known schedule and known probability distribution. This allows traders to calculate the fundamental value easily at any time.

As an example from one experiment by Ackert et al. (2006), each trader was given two shares of an asset that each period paid a dividend of \$0.50 with probability 48%, \$0.90 with probability 48% and \$1.20 with probability 4%. The expected value of the dividend in each period was:

$$0.48 * \$0.50 + 0.48 * \$0.90 + 0.04 * \$1.20 = \$0.72$$

As the experiment ran for 12 period, the expected value of the asset in period 1 was \$8.64. To calculate the expected value at any time, multiply \$0.72 by the number of remaining periods.

For a risk neutral trader, this expected value is the fundamental value of the asset. For a risk averse person, the fundamental value would be lower as they must be compensated for the risk.

In the Ackert and friends experiment there was five minutes of trading using a computerised platform before each dividend payment. A plot of the price in four markets over the 12 periods, plus the fundamental value, is shown below.

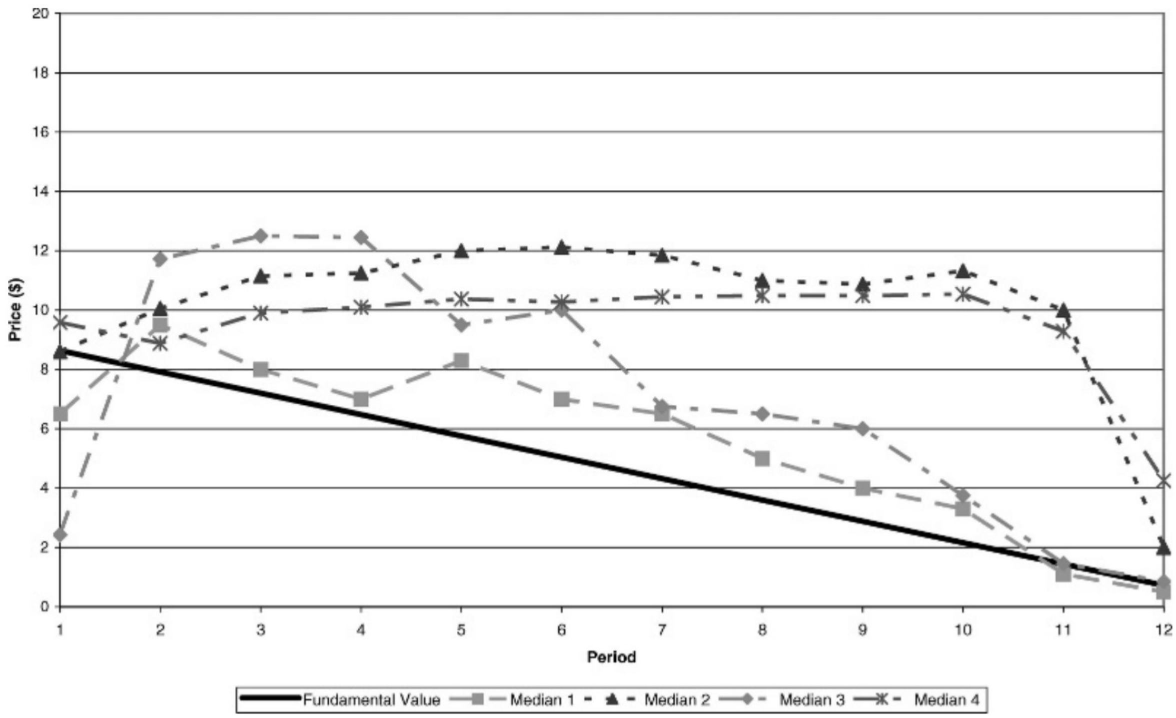

This price pattern is typical of that in experimental bubble markets. Although the price is below fundamental value in two of the markets in the first period, it quickly rises above the fundamental value. Often the bubbles are persistent - as are two here - before finally crashing just before the music stops.

Importantly, these bubble are created despite the intrinsic value of the security being trivial to calculate at any point of time. That, of course, is not the case in a real market. But the fact bubbles can occur when the value is obvious suggests that uncertainty cannot be a complete explanation for their emergence.

Although not explored in the papers noted above, one other themes that has emerged from these experiments is that with experienced traders or bubble market experimental participants, the bubbles are milder and disappear earlier.

9 Excessive volatility

If you look at a chart of stock prices, one striking feature is that they are volatile.

Under the efficient markets hypothesis, variation in stock prices should be due to new information. Shifts in prices should also be accompanied by changes in rationally-varying forecasts in dividends.

A challenge with testing whether changes in prices are caused by changes in expectations of future dividends is that we don't know what dividends will be paid in the future. However, Shiller (1981) noted that we could look at past dividends to construct the "ex-post rational stock price" for any point in time given dividends paid from that point in time until today.

You might note that even if we look at past prices and the dividend flows from then on, the expectation relates to the discounted sum of dividend flows through to infinity. We will always lack some information to test whether the price change is justified.

Shiller argued that the ex-post rational stock price should be more volatile than the price itself. This is because future dividends that comprise the ex-post rational stock price will move based on unexpected information - information that investors did not have when the price was set - adding volatility to the ex-post rational stock price that the price itself should not be subject to.

In the chart below, the solid line represents the real Standard and Poor's Composite Stock Price Index, detrended for long-run growth. The dashed line is the ex-post rational stock price. It is obvious that the volatility of the index is far above that of the ex-post rational stock price.

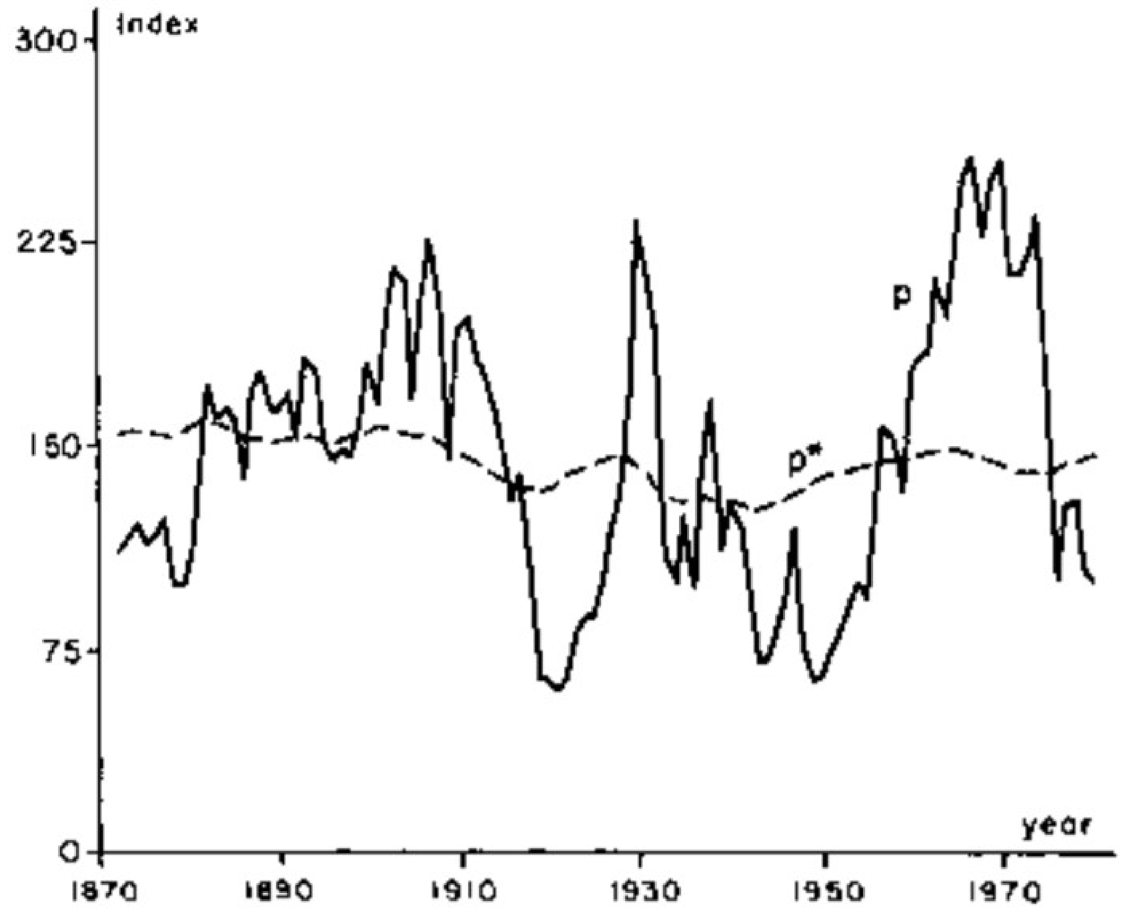

FIGURE 1

Note: Real Standard and Poor's Composite Stock Price Index (solid line p) and *ex post* rational price (dotted line p^*), 1871–1979, both detrended by dividing a long-run exponential growth factor. The variable p^* is the present value of actual subsequent real detrended dividends, subject to an assumption about the present value in 1979 of dividends thereafter. Data are from Data Set 1, Appendix.

9.1 Behavioural explanations

There are many explanations for excessive volatility (beyond the plentiful arguments that the volatility is not actually excessive). Many of these are based on similar ideas to the other anomalies that we have discussed.

One relates to representativeness. When an investor sees a short sample of good earnings, they believe that earnings have gone up and will continue to be higher in the future. They therefore over-react. They might also extrapolate past returns too far into the future when estimating future returns, again based on representativeness.

9.2 Further material

If you wish to read further material on behavioural finance, Shiller (2003) is one good starting point.

The following video, an interview with Eugene Fama and Richard Thaler, covers many of the behavioural finance topics discussed in this book.

References

Ackert, L. F., Charupat, N., Church, B. K., and Deaves, R. (2006). Margin, short selling, and lotteries in experimental asset markets. *Southern Economic Journal*, *73*(2), 419. https://doi.org/10.2307/20111899

Ball, R., and Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, *6*(2), 159–178. https://doi.org/10.2307/2490232

Barber, B. M., and Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The Journal of Finance*, *55*(2), 773–806. https://doi.org/10.1111/0022-1082.00226

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*(3), 307–343. https://doi.org/10.1016/S0304-405X(98)00027-0

Benartzi, S., and Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*, *110*(1), 73–92. https://doi.org/10.2307/2118511

De Bondt, W. F. M., and Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance*, *40*(3), 793–805. https://doi.org/10.1111/j.1540-6261.1985.tb05004.x

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, *25*(2), 383. https://doi.org/10.2307/2325486

Fama, E. F., Fisher, L., Jensen, M. C., and Roll, R. (1969). The adjustment of stock prices to new information. *International Economic Review*, *10*(1), 1–21. https://doi.org/10.2307/2525569

Fama, E. F., and French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, *47*(2), 427–465. https://doi.org/10.1111/j.1540-6261.1992.tb04398.x

Fink, J. (2021). A review of the post-earnings-announcement drift. *Journal of Behavioral and Experimental Finance*, *29*, 100446. https://doi.org/10.1016/j.jbef.2020.100446

Jegadeesh, N., and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*(1), 65–91. https://doi.org/10.1111/j.1540-6261.1993.tb04702.x

Lakonishok, J., Shleifer, A., and Vishny, R. W. (1994). Contrarian investment, extrapolation, and risk. *The Journal of Finance*, *49*(5), 1541–1578. https://doi.org/10.1111/j.1540-6261.1994.tb04772.x

Malkiel, B. G. (2005). Reflections on the efficient market hypothesis: 30 years later. *Financial Review*, *40*(1), 1–9. https://doi.org/10.1111/j.0732-8516.2005.00090.x

Mehra, R., and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, *15*(2), 145–161. https://doi.org/10.1016/0304-3932(85)90061-3

Rabin, M., and Thaler, R. H. (2001). Anomalies: Risk Aversion. *Journal of Economic Per-*

spectives, *15*(1), 219–232. https://doi.org/10.1257/jep.15.1.219

Rendleman, R. J., Jones, C. P., and Latané, H. A. (1982). Empirical anomalies based on unexpected earnings and the importance of risk adjustments. *Journal of Financial Economics*, *10*(3), 269–287. https://doi.org/10.1016/0304-405X(82)90003-4

Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia*, *98*, 108–113.

Shefrin, H., and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, *40*(3), 777–790. https://doi.org/10.1111/j.1540-6261.1985.tb05002.x

Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*, *17*(1), 83–104. https://doi.org/10.1257/089533003321164967

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review*, *71*(3), 421–436. https://www.jstor.org/stable/1802789

Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, *56*(5), 1119–1151. https://doi.org/10.2307/1911361