12

Perfect and Efficient Markets, and Classical and Behavioral Finance

How Trustworthy are Market Prices?

This chapter explains the concept of an efficient market, which is not as strict as but closely linked to that of a perfect market. A market is said to be efficient if it does not ignore available information. To illuminate perfect and efficient markets, this chapter also explains arbitrage, an essential concept of finance, without which no study of finance would be complete. We then discuss the consequences of the concepts: What do efficient and/or perfect markets mean for predicting stock performance? How should you interpret the success of famous investors (like Warren Buffett)? And how can you use the concept of efficient markets to run an event study to help assess the valuation impact of big corporate events?

12.1 Definition of Market Efficiency

A perfect market sets up stiff competition among many investors. This state of affairs forces them to use all available information as well as they possibly can. This is called market efficiency: a situation in which prices reflect *all* available information. In a fully efficient market, you should not be able to use any available information to predict future returns better than the market can. This is the Efficient Market Hypothesis (EMH). **Begin Important**

A price is called efficient if the market has set the price based on the use of *all* available information. (PS: It is not necessary that any one investor has all the information.) **End Important**

Warning: Market efficiency is a different concept from mean-variance efficiency (the efficient frontier), which is used in the context of portfolio optimization. Economists love the word “efficiency” and thus use it in too many contexts.2]subsubsect:eff-frontierMean-variance efficiency

Figure **??** illustrates the efficient-market concept. Suppose investors consider an expected rate of return of 10% on ABC stock to be a fair rate of return, given ABC's characteristics. This figure of 10% could come, for instance, from the CAPM. Market efficiency then pins down the relationship between the best estimate of the price next year and the price today. In our example, if the market expects ABC to trade for $55 next year, it should set the price today at $50. The market would not be efficient if it had set today's price at $49 or $51. You can turn this around, too. You should not be able to locate information that tells you today when/if/that the true expected value tomorrow is really $60 (for an expected rate of return of 20%) or $40 next year (for an expected rate of return of –20%). If you could find information telling you with good confidence that a better estimate of next year's price is $60 (or $40), then ABC's stock would be mispriced. A market that has overlooked your information would not be deemed efficient. (Your obtaining this information would, over time, become quite profitable for you.)

(Omitted fig)

The practical use of the "efficient markets" concept raises two questions:

1. Where does the figure of 10% come from? It has to come from some model that tells you what rate of return ABC should have to offer given its characteristics, such as risk, liquidity, and so on. The CAPM is such a model (though a rather miserable one). Without a good model of what you should expect the rate of return to be, market efficiency is too vague a concept to be meaningful.
2. If the market is not perfect and different investors have different information, then exactly what information set are we talking about? If you are ABC's CEO, then you may have more information than the public. You may know whether the government will open a fraud investigation against you and whether you have the next new hit drink in the lab right now. You could know whether $50 today is too high or too low. Put differently, the market may be efficient with respect to publicly available information, but it need not be efficient with respect to insider information.

What should you conclude if you can determine authoritatively that the expected rate of return is really 20%? (This can happen either if you determine that the expected payoff is $60, not $55, or if the expected payoff is $55, but today's price is $45.83.) You could now draw one of two conclusions:

1. The CAPM is not the correct model. Instead, the market followed some other pricing model and wanted to set the expected rate of return for ABC at 20% in the first place.
2. The stock market is not efficient.

Can you see why market efficiency is so difficult to prove or reject? If you wish to proclaim a belief in market efficiency, and if you then find empirically that prices are not what your model predicted, you would simply proclaim that it was your model for the appropriate expected returns in your financial market that was wrong, not that the market was inefficient. It was your fault, not the market's. You just have to go back and search more — possibly forever — until you find the right pricing model.

Short-Term versus Long-Term Market Efficiency

Over long horizons (say, 1 year or longer), market efficiency is extremely difficult to disprove. The reason is that no one knows exactly what the correct model of pricing is. As you saw in Chapters and , benchmarking stocks to peers is a reasonable method that does, but its empirical reliability is modest. (And the CAPM is empirically simply wrong, despite its great intuition). Simply put, pick any stock, say ABC. We are not sure and cannot agree whether ABC should earn 10%, 20%, or 30% a year. This ambiguity renders market efficiency a concept that itself in practice often evades empirical testing. It is also why market efficiency is sometimes (perhaps unfairly) derided as being more religion than science. Based on the existing long-run evidence, some reasonable analysts conclude that financial markets are generally efficient (and our [CAPM] pricing model is wrong); and other reasonable analysts conclude that financial markets are generally not efficient.

Of course, in extreme circumstances, market efficiency can be a useful claim even on such long horizons. We know that no reasonable model of financial markets should give investors great bets like "+$1 million with 99% probability and –$1 with 1% probability." Expected returns this high would be way out of line with *any* reasonable pricing model. Even expected rates of return of 100% per year would surely be unreasonable for (most) stocks. Of course, few people doubt that big U.S. financial markets are, to such a first approximation, efficient — we all know that you just can't earn that much. But there is a large gray zone where it is difficult to distinguish between model error and market inefficiency. Because no one knows for sure what the correct model of expected stock returns is, no one can tell you affirmatively whether the stock market set the price of ABC stock so as to offer investors an expected rate of return on ABC of, say, 10% a year or 12% a year.

However, over short horizons (say, a day or so), market efficiency is a surprisingly useful concept. The reason is that over a single day it does not matter as much whether you believe the expected rate of return on ABC is 0%, 10%, or 20% per annum. Even on the high end of 20% per annum, the expected rate of return is still only about 5 basis points per day. Roughly speaking, regardless of whether you believe in the CAPM or not, you should expect day-to-day returns to be just a tiny bit above 0%. You should attribute most daily price movements to random fluctuations, presumably caused by unpredictable news of changes in the economic environment. However, if you can predict day-to-day stock movements (and you have thousands of days of historical stock returns to work with), then chances are that you would not blame the pricing model. Instead, you would probably conclude that the market is not efficient.

Noise:
01-05

Begin Important

- Over short time intervals (say, days), market efficiency is a very powerful concept. The expected rate of return should be tiny. If it is different, the market is probably inefficient.
- Over long time intervals (say, months or years), it is difficult to pin down what the appropriate expected rate of return is. This makes it difficult to disentangle errors in the pricing model from market inefficiency.
- Prices should move only when there is news about future cash flows or discount rate changes, where news is defined as the unanticipated component of new information that is arriving. Such news can be firm-specific or market-wide.

End Important

Relation to Perfect Market

Although the efficient market concept is different from the perfect market concept, the two are intimately linked — in fact, so much so that they are often casually confused. The reason is that if a market is perfect, economic forces drive it instantly toward market efficiency. Put differently, if a market were perfect but inefficient, it would be too easy to get rich. Too many smart investors would trade and the inefficiency would disappear. Market prices would instantly adjust to prevent this from happening. Therefore, if a market is perfect, it is inevitably also efficient.

The converse is not true, however. It is quite possible for an imperfect market — for example, one in which there are taxes or different opinions — also to be efficient. You could even (crudely) think of market efficiency as the result of the trades of many investors with many different information sets (opinions). The market price is the outcome at which investors do not wish to trade further. Appropriately weighted by investment amounts, one half believes the market price is too low; the other half believes it is too high. Of course, efficiency should be contemplated market by market. It is probable that some financial markets are efficient while others are not. The closer a market is to being perfect, the more likely it is to be efficient.

Another way to understand the difference is to compare assumptions. Of the four perfect-market assumptions, only one has any overlap with and bearing on the efficient-market concept: the one regarding "same information set and opinions." And even the information requirements are weaker. (The other three assumptions may prevent exploiting it.) It is not necessary that all investors have the same information and opinion (as in the perfect-market setting), just that the market price is the same "as if" the market itself had access to all the information at once. So, a market can be efficient even when investors know different bits and pieces of information and/or have different opinions, just as long as the market-price is the same that it would be if they were all sharing their information and opinions.

Perhaps the most important perfect market assumption driving prices toward efficiency is the absence of transaction costs. Without them, it is easy for you and other investors to trade on any information that the market has not yet incorporated in the stock price — and thereby earn an unusually good expected rate of return. However, the no-free-lunch axiom applies here, too. High transaction costs would make it more likely that you could expect to find violations of efficient markets. But if it is very expensive to trade and if the market is therefore not efficient and does not respond to news immediately, it would also be very difficult for you to take advantage of such inefficiencies.

Here is a practical example of how any market inefficiency would disappear quickly in a perfect market: What would you do if you learned that the market always goes down on rainy days and up on sunny ones? It is unlikely that the average investor requires extra return to hold stocks on sunny days — and, even if the average investor does, it is enough for you if you are not among them. You would never buy stocks when the weather forecast predicts that rain is coming. Instead, you would only buy stocks when the weather forecast predicts that the sun will shine. Investors like yourself — and there are of course many such investors in perfect markets — would rapidly bid up the prices before the sun shone, so that the prices would no longer systematically go up on sunny days. The end result is that if markets are efficient, then you should not be able to earn abnormally good sunny-day returns — at least not this easily. In a reasonable world, to earn higher expected rates of return, you must be willing to take

on something that other investors are reluctant to take on — such as higher portfolio risk. Today's weather alone should not do it. (Interestingly, academics do disagree on whether the weather in New York City [where a large number of traders are located] has a small influence on stock returns. Some papers claim it does, so that the market is inefficient. Others dispute this assertion, claiming the historical correlation is spurious and disappears if the statistical tests are done correctly. All agree that the weather influence is small, however.)

Conversely, it is easier to believe that markets are *not* (or less) efficient if transaction costs are high. But even if the market is not perfect, market inefficiencies should still raise eyebrows. For example, let's say that the appropriate rate of return on ABC was still 10% and the price was still $50. Alas, after you have run a few regressions, say you learn that the expected future price is not really $55 but $51. (The true expected rate of return would thus not be 10%, but 2%.) In a perfect market, some investors may want to short some ABC and use the shorting proceeds to buy another stock. This may not be possible if the market is imperfect and the costs of going short are too high. However, this leaves the question of why investors who already own ABC shares would not want to sell them ASAP. They would not incur the shorting transaction costs and would avoid the then lower-than-appropriate rate of return. (Maybe they are asleep?! Or maybe even the non-short related plain selling transaction costs are too high?!) Their economic self-interested behavior adds to the pressure in driving markets toward efficient pricing, even in a market that is imperfect.

Begin Important

- If a market is perfect, market forces should drive it strongly and quickly toward efficiency.
- If a market is not perfect, market forces should still drive it toward efficiency but not as powerfully. Third-party traders with higher transaction costs may not be able to aid in the process.

End Important

Market Efficiency in Modern Financial Markets

The financial markets for U.S. Treasuries, large-country currencies, *large* publicly-traded stocks and funds, and many other assets, seem reasonably close to perfect and thus efficient. They are very competitive. There are thousands of buyers and sellers, even thousands of tax-exempt investors, and modest transaction costs.

For U.S. Treasuries and OECD currencies, it seems unlikely that major investors would have true inside information. It also seems difficult to imagine that you or I could systematically outsmart the prices in such markets. After all, thousands of other traders are at least as smart as you or I. Smart trading desks have thoroughly explored the correlations in the data and deployed billions of dollars to exploit even the smallest of unusual opportunities. They would flock to good bargains and avoid bad bargains (along with us). It would be astonishing if these markets were not very efficient.

However, for small firms, it is possible that the stock market may be less perfect and less efficient. Some insiders may try to hide their trades from the SEC. There are many small stocks (even on the NASDAQ exchange) that trade only rarely and have reasonably large transaction costs.

- The bid-ask spread is often high.
- The posted bid-ask spread is only guaranteed for 100 shares — if you want to trade more shares, the price is likely to move against you.
- Commissions can be high.
- Shorting small stocks can be very costly when compared to the ideal of a perfect world in which you have full access to the proceeds (e.g., to earn interest).

In a round-trip transaction, you will face the first three issues once when you buy and once when you sell. Thus, it is unlikely that small stocks will immediately and fully reflect all information appropriately. The historical prices you see posted may be "stale" and may not even reflect the price that would have applied if you had wanted to trade.

Market efficiency is never white or black, but always a shade of gray — just as it is for perfect markets. Large, liquid S&P 100 stocks are pretty close to efficient; small NASDAQ stocks may not be. One conceptual question that had vexed academics for a long time was how markets could be efficient to begin with. After all, if there is no money to be made, why would anyone bother collecting information on firms? And if no one bothers to collect information on them, how can the market incorporate all information and thus be efficient? Eventually, a resolution to this puzzle was offered by Grossman and Stiglitz. They argued that markets can never be 100% efficient — they can only be, say, "99%" efficient. In equilibrium, good information collectors should earn just about enough trading profits to break even on their costs of information collecting. On the margin, the expected costs of learning and trading on more information are exactly equal to the expected trading profits. The informed investors earn this money trading against noise trader[noise traders], who do not collect information and who may trade for idiosyncratic reasons (e.g., to pay for tuition).

(Omitted anecdote)

The fact that large-firm stock markets are pretty efficient means that, by and large, you can trust these financial markets to get asset values about right — at least within the limits of the typical transaction costs — and to get it right *immediately*. As an investor, would you not rather face an inefficient market? If it were inefficient, you might be able to find some good bets (opportunities that earn unusually high expected rates of return). But it would not all be gravy. In an inefficient market, you could not rely on market prices being fair — they could be inappropriately too high or too low. You would never really know whether you are overpaying or underpaying. Investing would be a very messy business. You might have to spend a lot of time and money to determine whether prices are fair. The advantage of efficient markets is that if you hold a portfolio of many large and liquid stocks, you do not have to spend a lot of time and money to perform due diligence in order to determine whether stocks are fairly priced. All you need to do is to make sure you are appropriately diversified to meet your risk-reward preference (and watch your taxes). And you can probably accomplish this goal by buying just a few large index-mimicking mutual funds — the most popular personal investment strategy among finance professors by far.

(Omitted solvenow)

12.2 Market Efficiency Classifications and Behavioral Finance

A firm belief in efficient markets is what defines a school of thought known as classical finance, an outgrowth of the school of rational economics. This belief assumes that the evidence supports the efficient market hypothesis, or EMH, which holds that all securities are priced efficiently. In contrast, another school of thought, often dubbed behavioral finance, posits that markets sometimes do *not* use all available information. Depending on how strong a believer in classical finance versus behavioral finance you are, you may believe that there are no especially good trading opportunities, few trading opportunities, or plenty of trading opportunities. Both camps agree, however, that market perfection plays a crucial role in determining whether a particular market is efficient or not.

Almost all financial economists, regardless of camp, believe in basic market efficiency for large markets and liquid securities. Not even behavioral economists believe that it is easy to get rich trading on easily available information. Instead, the disagreement between more classical and more behavioral economists is, loosely, about whether stock markets are "99% efficient" or "97% efficient." Classical finance believes in the former, behavioral finance in the latter.

Because you can trade millions of dollars in large-firm stocks or market indexes relatively easily and at low transaction costs, it may not require huge efficiency violations for behavioral-finance traders to earn money and for classical-finance traders to be left behind. Exploiting just the tiny — say, $100\% - 97\% = 3\%$ — violations from market efficiency could make you a star investor. (This is also not coincidentally why so many fund managers publicly proclaim their faith in behavioral finance.) However, don't take me too literally here — the 99% versus 97% is an analogy, and there is really a spectrum of beliefs in market efficiency among economists and fund managers. Now, although you should realize that any classification scheme really identifies just segments on a continuous line, you can still try to classify financial economists and investors by their relative faith in efficiency. Let's look at some such classifications.

The Traditional Classification

The traditional definition of market efficiency focuses on information. In the traditional classification, market efficiency comes in one of three primary degrees: weak, semi-strong, and strong.

Weak market efficiency says that all information in past prices is reflected in today's stock prices so that technical analysis (trading based solely on historical price patterns) cannot be used to beat the market. Put differently, the market is the best technical analyst.

Semistrong market efficiency says that all public information is reflected in today's stock prices, so that neither fundamental trading (based on underlying firm fundamentals, such as cash flows or discount rates) nor technical analysis can be used to beat the market. Put differently, the market is both the best technical and the best fundamental analyst.

Strong market efficiency says that all information, both public and private, is reflected in today's stock prices, so that nothing — not even private insider information — can be used to beat the market. Put differently, the market is the best analyst and cannot be beat.

In this traditional classification, all finance professors nowadays believe that most U.S. financial markets are *not* strong-form efficient: Insider trading may be illegal, but it works. However, there are still arguments as to which markets are only semi-strong-form efficient or even only weak-form efficient.

The Fundamentals-Based Classification and Behavioral Finance

I prefer an alternative classification of market efficiency, which divides economists based on their beliefs in whether prevailing market prices reflect underlying values:

A true believer would argue that financial prices always reflect the best net present value estimate of all future cash flows. This means that stock prices should change correctly if and only if news about fundamentals (cash flows or discount rates) appears.

A firm believer would argue that financial prices may sometimes deviate from the appropriate best estimate of future cash flows. However, transaction costs make it practically impossible for investors to find unusually good bets.

A mild believer would also argue that financial prices may sometimes deviate from the appropriate best estimate of future cash flows. However, unlike a firm believer, a mild believer would argue that there are occasions when it is possible to exploit this misvaluation. This would result in the occasional unusually good bet. Usually, the profitabilities of such bets should remain within economically reasonable magnitudes — a couple of percentage points a year on the high side. Mild believers thus think that smart fund managers can offer investors slightly better bets, but nothing more. There are no guarantees.

A nonbeliever would argue that financial prices regularly deviate from the appropriate value, and to an extent that allows investors to obtain great bets fairly routinely.

These classes are progressively weaker along the market efficiency dimension. For example, a firm believer need not be a true believer. Firm belief can be the right club to join if financial price changes are indeed unpredictable, but not because of news about fundamentals. There could be unrelated noise in stock price changes, especially in the short run. A mild believer need not be a firm believer: Transaction costs may be low enough to permit great trading strategies based on EMH violations. A nonbeliever need not be a mild believer: Financial markets may just beg to be exploited. This classification is related to but not the same as the earlier classification. For example, it is possible that markets do not reflect all fundamental information, yet stock returns are unpredictable.

Where do most finance professors sit in this classification of beliefs? Virtually no academic is a perpetual nonbeliever, and only a very few remain in the "true believer" camp. Instead, most finance professors are somewhere between the "mild believer" camp (the center of behavioral finance) and the "firm believer" camp (the center of classical finance). The debates between the two more extreme sides of these camps — the more "classical rational economists" and the more "behavioral economists" — remain interesting. After all, bringing new evidence to bear on these disagreements is the process by which we learn more.

My Own Assessment

Here is my own view. Be warned: smart people can come to different conclusions. As far as market prices are concerned, I am mostly in the firm-belief camp. As far as other price patterns, activities, or beliefs are concerned, I am more in the mild-belief camp.

I believe that ordinary individual investors are unlikely to be able to find investing strategies in the public financial markets that earn high excess returns. A very few sophisticated funds may be able to earn a few basis points extra per year. But these funds are scarce and I don't know how to identify them. Even after decades of research that has tried to identify them, academics have usually found only that about half of all funds outperform the market and half underperform the market — and this is *before* we subtract out fund transaction costs.1.4]subsubsect:empevidencefundsEvidence on Fund Performance Most finance professors stash their own money in low-cost passive index funds.

But I also believe that most investors are not necessarily rational. Pundits love to talk about "investor psychology" or "herd instincts." For example, investors think they are smarter than they are — that they can predict when specific stocks are overpriced and underpriced. (I know I have done it, too.) This is why I believe that trading in the stock market seems so (inexplicably) active. It is also why many investors remain woefully underdiversified. For example, Nobel-prizing winning research has shown that investors are "loss-averse." It seems very plausible that loss aversion influences their stock trading patterns. This can make them incur unnecessary capital gains taxes at the end of the year.

Nevertheless, I remain in the firm belief camp as far prices are concerned. I don't find it so plausible that, say, loss aversion necessarily would influence prices. I have two reasons. The first is that different investors would have started out at different investment levels. They would thus suffer from loss aversion relative to different starting points. In the aggregate, prices would not necessarily behave as if there was one loss-averse investor who bought shares at a particular price point. The second is that, if prices were badly set, a few smart investors would try to take advantage of this behavioral bias. They would quickly drive prices back to where they would become efficient again.

(Omitted solvenow)

12.3 Prominent Market Inefficiencies and Limits to Arbitrage

Occasionally, there seems to be evidence that seems to refute market efficiency. Let's discuss the most prominent ones.

Equity Carveouts

Famous finance professors Owen Lamont and Richard Thaler described a dramatic example of market inefficiency that occurred in 2000. The network company 3COM had spun off the personal organizer device company PALM (whose products were later wiped out by Blackberry devices, which in turn were wiped out by Apple iPhone and Android devices). Widely reported in the press at the time, 3COM retained 95% of PALM's stock — and announced

that each shareholder of 3COM would soon receive 1.525 shares of PALM. After IPO, PALM closed at $95.06 per share. Therefore, 3COM should have been worth at least $1.525 \cdot \$95.06 \approx \$145$. Instead, 3COM shares closed at $81.81 in March 2000. This was almost surely a violation for true believers in market efficiency.

However, three other professors Cherkes, Jones, and Spatt then explained why it was impossible to get rich from the discrepancy. Taking advantage of the discrepancy would have involved going long in 3COM and short in PALM. Unfortunately, it was practically impossible to find PALM shares to borrow (a requirement for shorting in the United States). PALM shares later "enjoyed" an almost uninterrupted fall in price, from $95 in 2000 down to less than $2 per share by 2003, thereby wiping out the paper opportunity. In sum (pun), even though this episode rejected "true belief in market efficiency," it did not reject "firm belief in market efficiency." The price discrepancy was not easily exploitable.

Similar violations of summing-up constraints have also occurred a few times in other stocks since then, usually with similar constellations of EMH deviations and transaction costs. Smart investors quickly try to exploit them which usually makes them disappear.

Factor Anomalies: Value and Momentum

Beginning in the 1990s, finance professors have been regularly publishing academic papers claiming that some new quantitative strategy or another has outperformed reasonable average rates of return historically. The most prominent ones were forms of value investing strategies (buying boring old-economy stocks, selling glamorous high-growth new-economy stocks) and momentum investing strategies (buying stocks that have gone up and selling stocks that have gone down over the 12 months). Academic papers suggested that these strategies would have offered "excess returns" as high as 1-2% per month. In 2013, researchers from one of the largest hedge funds in the world published a paper with the provocative title "Value and Momentum Everywhere." It became one of the most cited papers of the decade. What could possibly go wrong?

➤ Value Investing

Value investing originated after the stock-market crash of 1929 (the famous "Black Thursday" of Oct 24, 1929, and "Black Tuesday" of Oct 29, 1929). Benjamin Graham and David Dodd (two Columbia professors) published a famous book called Security Analysis that argued that it is better to invest in firms that are less flashy and more down-to-earth (having high book values, low market prices, high dividend yields, etc.). Their most prominent acolyte was himself the most famous investor of the 20th century, Warren Buffett. In the early 1990s, the two most famous finance professors of our era, Eugene Fama and Ken French, analyzed the stock return data more systematically and came to a similar conclusion. They defined value firms as those with high book-to-market ratios and growth firms as the opposite. "Value firms" had indeed outperformed "growth firms" for many decades. They called this better performance the value effect.

Interestingly, Fama and French did not view this empirical pattern as a rejection of market efficiency. Instead, they believed it was due to some as-yet-unidentified risk premium for holding value stocks presumed to be riskier than growth stocks. However, I found it difficult to see what the nature of this value risk was. If anything, value stocks seemed less risky, not more risky — they had lower standard deviations of return and lower market betas

than growth stocks. From the standpoint of viewing risk in terms of exposures to other risks, no one knew what kind of insurance these growth stocks were offering that made investors content to accept lower returns. Therefore I agreed at the time with the behavioral finance research that interpreted the value premium more as a market inefficiency problem. Whatever the cause for the value effect was, the result was that the investment strategies of most quantitative hedge funds over the last two decades have been heavily influenced by tilts towards value, as have many published papers by academic researchers.

]sect:many-asset-classesVanguard Value and Growth Fund One way to examine the value effect is to compare the performance of the Vanguard value fund (VVIAX) and the Vanguard growth fund (VWUAX). Over our 17 years from 2005-2021:

	Ari	Sdv	Geo	# Years Better
Vanguard Value	10.2%	16.1%	8.8%	6
Vanguard Growth	14.4%	21.1%	12.4%	11

Similarly, Ken French has posted the Fama-French HML factor for decades. This factor is a portfolio investing long in high-value firms and shorting high-growth firms. Its net returns look similar:

	Ari	Sdv	# Years Positive
Fama-French Value Minus Growth	–2.9%	16.7%	7

Ending in 2021, value did not outperform growth for investment-periods that have begun anywhere from 1997 to early 2020. Compared to growth stocks, value stocks have been miserable investments for at least 20 years. My view now is that whatever value investing may have been twenty years ago, it is probably gone as of 2020. Value stocks may outperform or underperform but it is now quite unpredictable.

➤ Momentum Investing

The second-most important factor anomaly was the momentum investing strategy. Momentum investing strategies require going long in stocks that have increased greatly over the last year, and going short in stocks that have decreased greatly. (It requires a few more contortions, but this is a reasonable characterization.) As with value, behavioral finance researchers were quick to adopt momentum as a consequence of investor psychology. They also developed plenty of theories that argued about how the psychology of investors could explain momentum.

Yet over the last 17 years, Ken French's data suggests that the average rate of return on the momentum investment strategy was — drumroll — 0.03% with a standard deviation of 23.8%. This rate of return is statistically and economically insignificant. Momentum investing also had the unpleasant aspect of sudden nasty risk. It lost 83 cents for every dollar invested in 2009!

➤ But, but, but ...

A scouring of the data — some would call it torture — can still find momentum and value effects in some other markets. One can juice up the returns by investing more in some years only. And both value and momentum strategies still look good over much longer time-periods. It's just that they have not performed particularly well for

about two decades. And presumably you don't care much whether value and momentum did well when the Beatles were at the top of the charts. You probably care more about whether they will do well in the 2020s and 2030s.

As for myself, I would not invest my own money into the thesis that value and momentum will or will not return. Frankly, I am not confident that there is *any* factor strategy that has a positive average rate of return *forward-looking*. This opinion is based on four observations. First, it is easy to find a piece in the hay stack that looks somewhat like a needle but is not. There is a lot of data and with enough analysis, it is easy to find some spurious correlations. Second, maybe there is some additional risk (like the 83 cent loss in 2009 for momentum.) Third, transaction costs in these strategies can be considerably higher than they are for buy-and-hold strategies. And fourth and perhaps most convincing, as two other academics (Pontiff-McLean) have pointed out, smart investors are reading the academic literature, too. When they try to take advantage of an anomaly, they also make it go away.

More than ever, I believe that the EMH has been winning this argument. It is very difficult to find stock selection strategies that can be expected to outperform the rate of return on the stock market. It is better to make sure to be well-diversified than it is to try to beat the market.

(Omitted solvenow)

Bubbles and Bitcoin

Many, but not all researchers — myself included — believe that corners of the financial markets are subject to bubbles. Most also believe — like me — that these are not easily exploitable, thus rattling our faith in true market efficiency but preserving our faith in firm market efficiency.

A bubble is a runaway market, in which rationality has (at least temporarily) disappeared. There is a lot of academic debate as to whether bubbles in the stock market have ever, in fact, occurred. A strong case can be made that technology stocks experienced a bubble from around 1998 to 2000. It is often called the dot-com bubble, the internet bubble, or simply the tech bubble. I know of good fundamental-based explanations as to why the NASDAQ Index climbed from 2,280 in March 1999 to 5,000 by March 2000 and why it dropped from 5,000 back to 1,640 by April 2001 — but no good non-bubble explanations for both.

Today, I am wondering whether crypto-currencies like Bitcoin are in a bubble. A Bitcoin is a specialized number sequence that satisfies certain mathematical properties. Finding them requires a lot of calculations, which in turn requires spending large amounts of energy. However, the sequences themselves are otherwise intrinsically worthless. Knowing this sequence will not get you cash flows in the future, as it would for a project in an NPV analysis. The NPV of all future cash flows is zero.

At this point, the proponents of Bitcoin usually argue that sovereign currencies and gold also have no use and cash flows down the line. What makes them valuable is only that people trust them. The "only trust remains" claim is false, because you can pay your tax obligations with sovereign currencies. The "no other use" claim is more correct, although gold can also be used for some other purposes, too — dentistry, jewelry, circuit boards. However, the cryptonites do have a bigger point that should not be lost: it is that people trust these currencies as stores of values — so why not also trust Bitcoin?

To understand why currencies are useful requires a quick trip into history. Before the invention of currency, when you had a goat and wanted to buy bread, you had to walk with your goat all the way to the market and find someone who needed a spare goat and happened to have spare bread. Thus, having a common currency as a medium of exchange was very useful. You no longer had to schlepp your goat all the way back home if there was no bread seller with the need for a goat at that very moment in the market. You could sell the goat for gold to one party and buy bread from another the following day.

So what is the intrinsic value proposition for Bitcoin today? We already have bank accounts and credit cards, so yet another alternative currency adds relatively little value (and this is even if Bitcoin even had the stability and ease of transactions that modern payment systems provide). It probably is also not a great advantage for most people that their Bitcoin can be stolen far more easily. Ordinary people don't like the idea of potentially losing all their money when some hacker breaks in and steals all their Bitcoin, and they, the victims, are left without any recourse to recover their loss. Better a bank that one can call up to complain.

This really leaves two viable uses for crypto currencies. The first use is for people who are not able to use the international world banking system and who need absolute secrecy with respect to their governments. I am not judging the legitimacy of other countries' laws, but Bitcoin allows criminals under these codes to evade local laws. Bitcoin is also often used as payment in ransomware cyber attacks.

The second use is for speculators who extrapolate the past (itself manipulated by impresarios and existing crypto holders) to the future. They are afraid that they have missed out on the greatest opportunity of a lifetime — the ability to earn huge amounts for nothing. It is simply greed. In this sense, crypto is like a chain letter or a Ponzi scheme. The last suckers to hop on will lose all their money. As I am writing this, Warren Buffett has just called Bitcoin rat poison. His partner, Charlie Munger, called it venereal disease. As a finance professor, my professional opinion is that Bitcoin is also not suitable either as rat poison or a venereal disease. But I do share the sentiment.

Short Squeezes and GameStop

In 2021, yet another new phenomenon emerged. Posters on Reddit's WallStreetBets forum managed to go viral and coordinate hundreds of thousands of investors to buy stock in an obscure company with an obsolete business model of brick-and-mortar sales of game cartridges with little future — Gamestop (GME). With each investor betting just a few dollars, the resulting horde had an effective risk aversion that was more akin to that of a billionaire. And they managed to coordinate collectively on a scheme that would have been illegal for an individual — a short squeeze.

In the United States, to short shares, you need to borrow them from someone else first. Now imagine I owned all the shares in XYZ and you did not know this. You would ask your broker to short XYZ. The broker would borrow XYZ shares from me, and you would then sell them in the market. I will proceed to purchase the stock you are selling. I again hold all the stock. But you have made a commitment to return the shares to me at my request. If I "call" my shares back in, where can you requisition them now? There really is only one party from which you can buy shares — me. If I charge you $1 million per share, you have to pay it. This is called a short squeeze. It is unclear whether intentional and coordinated short-squeezes are illegal. Even if they are, it would be unenforceable for widely dispersed and loosely coordinated short squeezes. It is also difficult to prove — trading is anonymous.

In 2020, GME traded for about $4/share. In 2021, after the horde was purchasing GME shares, its price increased all the way to $325/share. Not surprisingly, many short-sellers faced margin calls. They had apparently misspeculated. A number of hedge funds were among these short sellers. They had bet against the underlying business but eventually had to throw in the towel and repurchase shares at thirty times the price that they had sold them for just one month earlier. Some funds went out of business. Basically, they had mistakenly assumed that all they needed to do was to analyze the underlying GameStop cash flows (being true or firm believers in market efficiency). No one knew about this novel "horde risk."

When the stock briefly collapsed again, the losers were not just the hedge funds, but also many retail investors who had bought in later at the direction of the horde. However, many of these small investors had purchased less than one share. Many hundreds of thousands of investors may have lost grand sums of, say, $30 — in exchange for the entertainment value of having participated in the bubble of a lifetime. (GameStop still seems to be periodically affected by horde activity.)

Such coordinated "bank-run type" events have now made it more risky to bring stock-market valuation of individual securities back in line with their fundamentals. Nevertheless, it seems unlikely that it could greatly affect large and deep financial markets, such as those for VFIAX (S&P 500). As I wrote earlier, market efficiency is never perfectly white or perfectly black. Large markets are probably very efficient, while small markets may or may not be.

12.4 The Random Walk and the Signal-to-Noise Ratio

How does the EMH matter to you if you are an investor? In an efficient market, there should be no obvious signals to outperform the risk-adjusted appropriate expected return. Assets should already be fairly priced.

Yet even today, some still argue about market efficiency. Why is the debate so difficult to settle? It is due to the fact that the signal-to-noise ratio in financial returns is low. The signal-to-noise ratio analogy comes from physics — the signal (here the expected price change) is small compared to the noise (here the day-to-day price volatility).

Let me illustrate what we can work with. What are typical price change magnitudes? For example, Friday, February 18, 2022 was an ordinary day, though marked by some uncertainty about a potential invasion of Ukraine by Russia. Some stocks declined. Roku (ROKU) lost 23%, Draftkings (DKNG) 22%, Paramount (VIAC 18%, etc. Other stocks gained. Theralink (OBMP gained 80%, Amplitude (AMPL) 21%. The trading volume leaders were Palantir (PLTR) with 138 million shares and AMD with 115 million shares. 287 stocks hit a new 1-month high, 780 a new 1-month low. More stocks declined than gained. The S&P index declined by 0.72%, VFIAX by 0.70%. Its consumer staples gained a few basis points, its IT stocks fell by 1%, INTC losing 5%, Nvidia (NVDA) 3.5%, Apple (AAPL) 1.5%, and so on. Treasuries rose a little but remained largely unimpressed. The yield on 10-year Treasuries changed from 1.97% to 1.92%. The 1-month Treasury dropped from 6 bps to 3 bps.

Is this information useful for investing purposes? How should you interpret it? Read on.

A Random Walk

In a perfect market, the best way to think of investment asset processes is that prices should follow a random walk.* The formula for a random walk is

(Omitted eq)

Ignore the tiny drift m for a moment. The ρ coefficient on the price today has to be very close to 1.0. If it were different, you could easily make money. For example, if ρ were 1.01, the price would be expected to be 1% higher tomorrow. Buying the shares every day over 252 trading days, you would earn over 1,000% per annum. If ρ were 0.99, you would short the shares. Same 1,000% return. Ergo, the ρ coefficient has to be very, very close to 1.0, with only the m component pushing on the price *on average*.

This random walk can also be expressed in terms of a rate of return from today to tomorrow. With a little algebra,

(Omitted eq)

which says that prices are mostly unpredictable, though pushed around by noise with some standard deviation s and still growing by just the tiny bit of m.

Begin Important

In the financial market context, "random walk" refers to a process in which the *expected* value of investments tomorrow is (almost) the same as the value today. Technically,

(Omitted eq)

where m is a very small positive drift. (m is the subject of asset-pricing models.) **End Important**

Let's make it more concrete with an example (with realistic values) for illustration:

The "Noise" (s): What would be a good value for the standard deviation of the error term (s)? As you already know, standard deviations are reasonably stable and easy to estimate. You also already know about typical standard deviations in the U.S. stock-market.

The typical day-to-day standard deviation for individual U.S. stocks has been around 2-3% per trading day — of course, somewhat dependent on the specific stock. For well-diversified portfolios, like our stock-market portfolio index VFIAX, the standard deviation has been lower — perhaps 1-2% per trading day. In terms of annualized volatility, this translates into about $\sqrt{252} \times 2.5\% \approx 40\%$ per annum for individual stocks and $\sqrt{252} \times 1.5\% \approx 20\%$ for well-diversified market-type portfolios.

Incidentally, with its loss of 0.72 bps for VFIAX, February 18, 2022 was pretty much an ordinary day in terms of day-to-day volatility.

*Some folks who want to show off their math training use the term martingale instead of random walk. (It came from now-obscure 18th-Century gambling strategies.) For our purposes, the two mean the same thing.

The "Signal" (m): What would be a good value for the drift m? As you already know, expected rates of return for stocks (such as the equity premium) are notoriously difficult to estimate. (It is also why models such as the CAPM are struggling.) Rather than rehashing our earlier hand-wringing, let's just go with a large number for the sake of illustration: 13-14% per annum. Earning 13-14% per annum translates into an average return of about 0.06% (6 bps) per trading day.

Put this all together, and our illustrative portfolio process from today to tomorrow is

(Omitted eq)

Over a full year, for 252 trading days, we would expect to see

(Omitted eq)

(Omitted anecdote)

Empirical Evidence About Random Walks

Stock prices indeed tend to follow roughly a random walk. This means that you cannot easily expect to get rich by trading based on past prices. Let's look at two largely futile attempts to do so anyway — technical analysis and active investment fund management.

➤ Weak-Form Efficiency and Technical Analysis

The main point of the traditional classification of market efficiency — specifically, the "weak" version — is the claim that you should not become rich by trading a strategy that relies only on historical prices. The attempt to do precisely this is called Technical Analysis. There are even dozens of programs on the Internet that help you divine the future from past price patterns. The names of some patterns are reminiscent of astrology ("ascending triangle"). Other patterns are cloaked with more pseudo-scientific humbug names ("fast stochastic").

Let me start with some trick questions. Look at the three plots in Figure **??**. Do you think they show stock-market patterns? Perhaps. Does it make sense to think that all these patterns can predict the future? Absolutely not! In fact, these kinds of charts are as easy to read as tea leaves.

(Omitted fig)

Plot (a) displays perhaps too strong a regular cycling pattern. If this pattern indicated future returns, we could quickly become a wealthy technical analyst. We would buy the stock only when it has "bottomed out" — a pattern that you could reasonably detect if you saw a few days of upward trend. Now, if you look hard enough at the data, can you find some stocks in the real world that have historically behaved like plot (a)? Yes — because with 5,000 stocks and more than 5,000 mutual funds currently trading, by pure chance, maybe one or two could show a pattern that would look remarkably similar to this kind of cycle pattern. But, despite assurances from some stock analysts that you could have made money if you had just trusted their past cycle patterns (and that you should now let them invest your money for a fee), the patterns would *not* represent the future — they would just have been historical coincidence.

What about plot (b)? Is it also cycling too regularly? No! This one was real. It was the Vanguard energy stock portfolio (VGELX). Do I believe that this past pattern will allow me to predict its future? Not as much as I believe that the sacred chicken of Rome will. (The ancient Romans liked to read the appetites of sacred chicken — perhaps as good a prediction method as any when perfect markets are at play.)

Plot (c) looks a little too regular in sloping up. Nevertheless, it turns out that this was also a real pattern. It was the overall stock market, specifically VFIAX (S&P 500). Stocks were a remarkably good and steady investment (with a tiny positive m) over the last 17 years, with the exception of the Great Recession of 2008-9. But do not imagine that this tells you that any crashes will quickly be reversed in the future. (Incidentally, plot (a) just added a cycle pattern to the VFIAX.)

(Omitted fig)

An easier way to see the absence of any daily patterns — even in VFIAX — is to plot today's return based on yesterday's return. Figure **??** does just this. Can you see a pattern? I can't. Back to my sacred chicken. There is not much financial juice in trying to predict how a stock will perform tomorrow, given how it performed today. Similar conclusions apply if you extend your use of historical price information beyond yesterday.

You can even try out your own technical analysis at a number of financial websites, such as YAHOO! FINANCE: Look up any stock and choose "Charts," then "Technical Analysis." It is fun, but unfortunately as useless as learning to divine the sacred chicken of ancient Rome (https://www.anecdotesfromantiquity.com). Hey, at least the chicken helped build an empire — and I am not talking KFC. Shouldn't their historical performance count for something?

12.5 Your Great New Investment Signal

The question we now want to ask is what kind of data we would need to scientifically and statistically conclude that we have discovered an investment signal that "works" — i.e., that allows us to earn return above and beyond what the standard market random walk delivers.

Say our signal-based strategy could deliver an average of 2-3% extra per year. How hard could delivering 2-3% extra per year be? Very hard! Trust me that if you could repeat this high a performance above a benchmark reliably year after year, you would have placed into the rare category of investment superstars. Of course, it is common to randomly earn more than 2-3% for a few years. (Even with my own miserable investing luck, I have managed to achieve this result in some years!) But it is uncommon to do so systematically and reliably year after year after year.

Yet our own problem is different now: We need to answer the question of how we determine whether our presumptive signal delivered dumb luck or actual information.

An excess performance of 2-3% per year means a superior performance of about 1 bp per trading day. The problem is that we need to confirm the signal's validity in an environment where the noise is about 20-30% per year (or 150 bps per day). Incidentally, you could also view the noise as the "good luck" or "bad luck" aspect in this context.

With the noise (standard deviation) of about 100-200 times as high as the signal (expected rate of return), our signal-to-noise ratio is quite low. And if our signal gives us pointers only for a few stocks and projects here and

there, rather than for large diversified projects and portfolios all the time, then the applicable portfolio noise could not just be 20-30% but 30-50% per year instead. This would, in effect, halve our signal-to-noise ratio.

Our problem is to determine whether our signal is real or illusory. This means we must be able to distinguish between the usual 5 basis points and an unusual 6 basis points average daily rate of return in a soup of 150 basis points of noise every day.

Obviously, one daily observation is not going to cut the mustard. That is, if our signal-based strategy happened to earn 50 basis points today, we could not reliably conclude that it was our signal. In fact, if anything, we should believe that such a large return — 50 times more than our expected excess average return — was primarily noise.

1.96 in a stats course.

Now trust me (or recall from your statistics course) that the T-statistic is defined as the mean divided by the standard deviation, $\mathsf{E}(r)/\mathsf{SD}(r)$; and we "sort of trust" numbers with an absolute T-statistic of about 2.0 or more. If our strategy performs as expected, we expect a 1-day T-statistic of about 1 bps/150 bps ≈ 0.007. With only one day of data, our expected 0.007 is a long way off. We need more days.

How many? Let's ignore compounding and pretend that rates of return over time are just the simple sum of daily rates of return. In this case, our expected rate of return over D days is D times the expected rate of return over 1 day. Recall from Section]subsubsect:riskovertimeHow risk and reward grow over time that the standard deviation of our rate of return over D days is $\sqrt{D}$ times the standard deviation over 1 day. Our expected T-statistic over D days to detect superior performance is therefore

(Omitted eq)

If we expect a 1-day T-statistic of 0.007, we would expect a 100-day T-statistic of $\sqrt{100} \cdot 0.007 \approx 0.07$ — not even close to 2.0. We can estimate that we will need about 90,000 trading days to expect to reach a T-statistic of about 2.0. That's only about 357 years. And this works only if we expect our signal and the world still to work the same way over the next 357 years. Frankly, it's hopeless. We cannot possibly expect to validate this small a signal in this much noise within a reasonable time span.

Begin Important

- The quality of your inference about a strategy's performance increases roughly with the square root of time.
- On an average day, the typical stock may easily move up or down by about 50-100 times as much as it offers in expected rate of return.
- Therefore, it takes at least many decades, if not centuries, of data to reliably conclude whether an investment strategy's performance is real or illusory. By this time, the process or the signal itself may have changed.

End Important

Of course, if our signal were ten times as good (delivering not 2-3% but 20-30% per year), then we could expect to cut the mustard with only about 35 years of data. However, who has a signal that can deliver 20-30% per annum?

In the competitive stock markets of today, where smart traders immediately flock to any halfway reasonable opportunity, this seems absurd.

Could you gain more observations if you use individual stocks instead of individual days to test your theory? Unfortunately not. Stock returns at the same moment in time are not independent realizations. They tend to move together. For example, if your signal suggested oil was going up, and all 100 oil stocks increased, you did not have 100 independent observations confirming your signal's ability to predict. You really only had one.

We have also ignored transaction costs so far. In an imperfect market with transaction costs, you can view the EMH in one of two ways:

1. It should hold if you work with post-transaction cost rates of return.
2. It should hold if reasonably many investors have very low transaction costs.

   There is a special case here. Investors who were already planning to buy the stock anyway could just delay until the signal turns positive. Investors who were already holding the stock and planned to sell it could just delay until the signal turned negative. In this case, they incur no extra transaction costs.

   However, if these specific investors are *all* asleep at the switch, it would be impossible for other smart investors to take advantage of their failures.

So the EMH need not hold perfectly in an imperfect market, but it should still be a fairly reasonable description of reality — at least within transaction cost bounds.

A warning: It may be difficult to earn 1 bp extra per day, but it is easy to waste 1 bp per day on transaction costs even though trading costs in equity markets are low — for small transactions, as low as 1 cent per share now. This means that 1 bp is the minimum spread on a $100 stock. Unfortunately, even lowly 1 bps accumulate. If your strategy requires daily trading, it means that you are already 2-3% per year behind and your strategy needs to earn not 1 bp/day but 2 bp/day. If the signal turns out to be illusory, instead of being a superstar trader beating the market benchmark by 2% per annum, you would be lagging it by 2-3% per year.

(Omitted solvenow)

Investment Manager Performance Evaluation

Most smart professional fund managers have come to absorb these hard facts of life and developed proper humility. They focus on good diversification, tax minimization, pay-in and payout patterns, etc., sometimes tailored to the clients, and less on attempting to predict better or worse than the market.

(Omitted anecdote)

However, this is not necessarily the case for many analysts and talking heads. They want to make a mark for themselves. This is easier to do with extreme and outrageous statements. They could get famous if they turn out lucky. If not, what's the loss?

Analysts regularly claim to be able to identify stocks that are underpriced by 5%, 10% or more. For example, on Feb 17, 2022, Tesla was trading for $876. The 29 analysts following Tesla had an average 12-month price target on

Tesla of $1,121 (with a range from $300 to $1,580), suggesting an expected rate of return of 28%. They must indeed be brillant! While I worry about 2-3% per year, even subtracting out an estimate of overall stock-market rates of return, they purport to be able to predict an enormous rate of return.

If their signals are really this good, why are they even still on TV? They should be running their own funds and keeping silent about their financial prowess. Fortunately, because they have not, academic researchers were able to investigate their performance. Unfortunately, the empirical evidence suggests that analysts have been no better at predicting outcomes than dart throwers. So why do people still pay attention to these talking heads? I don't know. I am not a psychologist. Maybe "hope never dies."

But let's say we see someone who has invested really well — like a famous investor (e.g., Warren Buffett). How should we view Warren? Should we trust an "expert" with a glowing track record? Maybe. But let's consider the caveats.

First, recall that the low signal-to-noise ratio means it is difficult to determine why a particular trading strategy has earned high returns:

- Was the outcome due to sheer luck, which will not repeat (randomness)?
- Was it because it took on some risk that your appropriate return model forgot (your fault in measuring performance)?
- Or was it because the market was inefficient (you have a good signal, skill, and trading ability)?

on Hard Work, Luck, and
s: 2012-12-30

This is not just a problem for academics. In fact, we finance professors are lucky: We have continued to write papers that argue one side or the other even when we are skeptical about the future. Our money and jobs are not on the line. The real conundrum is faced by every investor in the real world every day. We have already talked about the challenge in some detail. How do you distinguish between a good and a bad signal — between skill and luck — when it comes to investing on your own or to selecting a fund manager?

But the problems go further. If you believe that the market is inefficient so that your investment manager can make you money, consider the following:

Is 3 years enough data? Recall our earlier conclusion that a strategy with great performance requires many decades before you can realistically conclude that it has worked. (This is assuming that the2]pg:howlongto-bereliable?Ascertaining superior performance world is not changing.) Few strategies have such long track records.

Remarkably, the most common industry standard for evaluating funds is their most recent three years of investment performance. There is no disagreement that most of the 3-year performance of funds is noise. This means that many investors (and especially investors in hedge funds) shift their holdings often based on noise. Why? Either they do not understand how long it takes to determine reliably whether a strategy works (possible), or they do not care too much about reliability (more likely). If they believe that there are many other strategies that also have a close to 50-50 probability of success, then eliminating one strategy that had 3 bad years and therefore only a 49-51 probability of success may not be a costly choice.

Are we sure about the risk? Here is another lesson for the wise (and unwise). Until 2008, I would have sworn that investing in momentum stocks was a strategy that was reasonably well-diversified and yet outperformed

the overall stock market. On average, it had delivered abnormal returns to the tune of about 5-10% per year. Stocks that have gone up over the last year and that are therefore momentum purchase candidates did not seem to be particularly risky. More importantly, momentum stock portfolios appeared well-diversified — a fact that should have moderated their ups and downs. Yet, after many decades of superior performance, in 2009, this momentum strategy suddenly lost 83 percent of its investment! (One plausible reason is that too many hedge funds were trying to chase momentum returns, and they all had to unload at the same time.) Which other seemingly great investment strategies are exposed to some risk that has just not shown itself yet?3]pg:peso"Peso" risk.

Lucky monkeys on keyboards? There are about 10,000 mutual funds]subsect:ADRsFunds-CAPMchapterMutual funds today that invest money on their investors' behalf. How many of them are likely to outperform the overall stock market next year (at least before they collect fees) if none of them has any superior investing ability? About 5,000. How many of these outperform the year thereafter? About 2,500. Even if there is absolutely no ability, pure randomness means that about 10 funds outperform the market every year for 10 years in a row. With enough candidates, some funds will inevitably produce consistently positive long-run track records.

Who is still here? What happens to the funds that have underperformed several years in a row? They disappear quietly. (In fact, they don't even need to appear. The SEC even allows a fund family to "incubate" funds privately for the purpose of obtaining track records. Start 1,024 of these funds, and after 10 years, you should expect to be able to go public with one of them that has outperformed 10 years in a row!) What happens to the funds that have outperformed several years in a row? They proudly announce their performances, advertise, boast, and collect more investments from outside investors. Their managers are supported by larger "research teams," appear better dressed and more "professional," and fly in executive jets. They are the ones who are most visible. Indeed, if you made money 10 years in a row in the stock market, would you not yourself believe that you have the ability to pick stocks?

Now put yourself in the shoes of an investor looking at the universe of mutual funds offered today. First, you won't notice funds that have performed poorly. They have already disappeared. Second, you will notice that the larger funds seem to have done better. On average, it will seem that currently available funds indeed can make you money — even if there were none in the world that could pick any better than chance. This phenomenon is called survivorship bias, because it means that you cannot consider the historical performance of existing funds to be a fair projection of their future performance.

Would we even see the evidence? Of course, maybe there are some investors who *can* pick stocks. Unfortunately, they would not want anyone to learn how they do it. In fact, they may want to do so secretly and privately, never eager to appear on anyone's radar screen. This opaqueness can make it difficult to find investors with superior ability and thus impossible to confirm their abilities.

Who would get the rents from trading ability? Even if the financial markets were inefficient and even if some fund managers could in fact systematically outperform the market, in a reasonable market, these fund managers would charge appropriately high fees to capture all the advantages that they provide to investors. After all, it is the fund manager who would have the scarce skill (picking stocks) and not the typical investor.

> Investors with money would compete to place money with such managers and accept higher and higher fund fees. In the end, it would be highly unlikely that uninformed investors could earn excess returns by investing in some manager's actively trading fund.

In sum, if you are looking for future performance, past performance may be your best guide. But always remember that recent past performance is still a very poor guide.

It is not an easy task to choose the right investments or investment manager. The best business to be in seems to be that of the investment manager. As a manager, you earn fees regardless of whether you make clients extra money or not. (In fairness, there are many good things that investment advisors can do, but selecting stocks to beat the market isn't one of them.)

Would it be better to have investment managers participate in the upside (as is the case for hedge funds), so that they have skin in the game? Maybe, but consider this: I give you stock tips, and I ask for money only if you make money. In fact, I only want 10% of your winnings. "You have nothing to lose." I only get something if I help you make money. Sounds like a deal? Now, if I pick a stock randomly, I have a 50-50 chance of making money. If you gain, I get something. If you lose, I pay nothing. In effect, I am arbitraging you! Remember, next time someone gives you a great stock tip, regard it with some skepticism: It probably has a 50-50 chance of being right. (Maybe I should give you the advice to buy a stock, and your neighbor the advice to sell it. This way, I will surely make money from one of you.) My only mistake is that I have told you my plan.

on Management Books:
6-30

➤ The Empirical Evidence for Active Management

So what is the empirical evidence? In general, it suggests that fund managers' luck is far more important than their ability. Whenever academics (or the *Wall Street Journal*) have searched for better performance among analysts or professional fund managers who have outperformed in the past, they have found little or no exceptional forward-looking performance. Table **??** shows a typical result in the literature: There were more funds that performed miserably than what we even would have expected by pure chance. Fewer than half of the funds could beat the zero benchmark. And many fewer funds than expected by random chance did great.

(Omitted tbl)

But what about persistence? Maybe there are some funds that are better than others? True. But the empirical evidence is again disappointing. Only about 54% of mutual funds that have outperformed their benchmarks over the last 1-3 years tend to outperform their benchmarks over the following 1-3 years. This is better than 50%, but not by much. And if you subtract fund fees, the average performance drops significantly below 50%. As fund prospectuses so aptly note — and as the empirical evidence so amply suggests — *past performance is no predictor of future performance*.

on past performance +
unds: 2013-04-17

There is a whole industry full of fund managers whose job it is to allocate assets to the actual investing funds. Chances are that your corporate pension fund will be managed by some. (So is mine. So are almost all university and non-profit endowments.) Of course, they all swear that they understand the problems but are immune to

* **Nerdnote:** There are some high-tech statistical techniques to take into account that researchers have searched, individually and collectively, for anomalies. This is beyond our scope.

them. They are professionals who *know* which funds are better than others. Could they *really* tell? Two finance researchers, Amit Goyal and Sunil Wahal, looked at 3,400 retirement plan sponsors from 1994 to 2003 and found that they were not particularly prophetic:

	Years Relative to Hiring	
	–2 to 0	0 to +2
Fired Funds	–1.6%	+3.1%
Hired Funds	7.6%	+2.3%

The pension funds hired fund managers *after* they performed well, not *before* they performed well. They fired funds *after* they performed poorly, not *before* they performed poorly. So why do these pension fund managers pretend that they can do a good job managing your money? Well, how much would you, as a client, be willing to pay for a plan sponsor who admitted to an inability to pick investment funds better than the sacred chicken?

There are, of course, other ways to make money: Warren Buffett's fund, *Berkshire Hathaway*, for example, runs many businesses, too. These businesses — like Berkshire's insurance and energy businesses — make money. But it is money earned the old-fashioned way — through hard work, liquidity provision, and risk-taking. .0]subsect:val-ueofliquidityBusiness of liquidity provision Writing insurance is risky business, and it deserves extra return. Warren Buffett himself is smart enough to acknowledge that the EMH is the most natural benchmark. He is on record as stating that "the professors who taught efficient market theory said that someone throwing darts at the stock tables could select stock portfolios having prospects just as good as ones selected by the brightest, most hard-working securities analyst. Observing correctly that the market was frequently efficient, they went on to conclude incorrectly that it was always efficient." Even Buffett is still a mild believer — he recommends index funds for most investors! Then again, he is paid not to be the firmest of believers.

In sum, most finance professors nowadays would agree that when one particular investor earns an unusual amount of money, even over a few years, it is usually more likely due to luck than to ability. The burden of proof is with the side that is claiming superior signals and investing ability — and a number of former finance professors have taken up the challenge and started their own funds. So far, none has done great. On the client side, if I were you, I would be very cautious investing my money, especially with any fund that charges high fees. Most finance professors invest their own money into low-cost index funds instead.

Begin Important

- Even in an efficient market, in which no one can pick stocks better than anybody else, with a very large number of investors, many will beat the market. A small number of investors will beat the market again and again.
- In the real world, there is little evidence that investors who did well picking stocks in the past are better at picking stocks in the future when compared to investors who did poorly.

End Important

(Omitted solvenow)

More Warnings

When fund managers earn great returns, they often become famous. To attract new investors, they then go out and talk more about their performance. The first targets are easiest to find at cocktail parties (you!) and industry conferences, itself an interesting business. Thereafter, it is usually admiring students looking for jobs. Enjoy these speakers and venues, but please remain skeptical.

A few lucky investors even go on to write books. Please read some of them. They all seem so sensible. All you need to do is to buy low and to sell high. Having sat through many presentations and having read many books, I can confidently state that about half emphasize the "buy low" while the other half emphasize the "sell high." When I am in a good mood, I can fake admiration for their "brilliant" investment insights. When I am in a bad mood, I offer somewhat cynical ambiguous praise that usually amuses only myself — such as "we would be lucky to get them to work with us." (These are our university donors, after all.)

A recent craze has focused on genetic algorithms and artificial intelligence. Unfortunately, even though they appear to be rather sexy high-tech sophisticated ways to pick investment assets, most lack something more basic: human intelligence. Fortunately, Michael Marcovici has remedied the situation with his "I Trained Rats to Trade, and Win, on Wall Street" true laboratory experiments. He taught lab rats to trade in the foreign-exchange and commodity-futures markets. Whenever they listened to a sound, the rats had to press either a green or a red button — green if they expected the prices to rise, red if they expected the prices to drop. The rats managed to outperform some of the world's leading human fund managers. Sure enough, a number of the rats outperformed some of the world's leading human fund managers. Unfortunately, the rats' ability to present their superior ability to potential investors (and thus generate higher fees) was limited by their lack of eloquence. (He should have gone with clucking chicken instead!)

Most funds write monthly communiques to their investors. They are largely collections of ex-post rationalizations and platitudes. There are some patterns. Funds on the up often write about the credit they deserve for their masterful insights. They describe competitive advantages, signals, edges, exciting and smart strategies, sentiment-reading abilities, contrarian acumen, etc. Funds on the down often write that nobody could have foreseen the problems — even Buffett read the tea leaves wrong, too; unprecedented market turmoil; irrational herd sentiments; unpredictability; temporary profit-taking; dollar averaging; the market failing to understand fundamentals; deteriorating data and decision making of others; fat-tail risk; dislocations; short-sellers; the Chinese, Russians, Saudis, Jews, or Arabs; the Fed doing too little or not enough, and so on and on.

Yet, the hard fact of life is that neither the ups or the downs are on target. Most hedge-fund performance in financial markets is just luck. In 2016, about a thousand hedge funds closed shop, mostly because of poor performance. There are a lot of one-hit wonders among them (and, of course, about half as many two-hit wonders, and about a quarter as many four-hit wonders).

Perhaps I shouldn't be so hard on the industry. We academics are really not so different. John Oliver's May 2016 show on Scientific Studies explains our shortcomings better than I can. I will try it anyway. If you read academic journals, you will find hundreds of papers showing how to beat the market. Just like fund managers, academics do not get rewarded for writing papers that opine that "the markets are fairly priced." They get rewarded for writing

papers that find that factor X had amazing returns. It's even better when the factor can claim to be related to behavioral and investor psychology — hedge funds and investors (also often our consulting clients) love such stories.

The problem is that even if each individual economist is (or were) scrupulously honest, as a collective, with thousands of us mining the data, we find many factors that seem statistically significant, yet are entirely spurious. Most of the time, quantitative hedge funds try to replicate and further test the academic factors right after the first public academic posting. If they confirm the findings, they then start trying to exploit past patterns. They all pile up into the factor as they back-test it, itself contributing briefly to some further good performance. (Virtually every academic equity fund has played "value" and "momentum" in some strategy or another, perhaps the reason for the disappearance of these effects.) And then, one day, the funds realized that they may have overreached.

What about me? Am I not brilliant? I placed large short bets on oil in 2013 when it traded above $100/bl. I believed long-term supply and demand could not sustain such a high price. In 2014, the oil price dropped below $50/b. This made my oil bets my best bets ever. I had talked about this in 2013 to my colleagues, who were then admiring my foresight. It was easy for me to rationalize how smart and prescient I was. But in all truth, my rationalization was really all non-sense. It really was primarily a gamble. I placed a bet, pure and simple. Ex-ante, other people just as smart on the other side believed the opposite. I happened to win. In financial markets, it is easy to place bets and someone ends up winning. In this case, it was me. Does this make me a genius investor? Or just a lucky one? In truth, with some reflection, probably more a lucky than a genius one.

Pennies Before Steamrollers

Sadly, we are not yet done. There is one more important point. Funds can wittingly or unwittingly give the illusion that they are better than random gamblers. It turns out that it is not difficult to show good historic performance *on average* and/or for a long time.

You can even show such good historical performance on a roulette table. It's called the "double-up strategy." For example, place money on red. If you lose (black comes up), place a bet twice as high. Do so until you win. Then go home and record today's investment performance as a gain. With a lot of money, it is likely that you will have years of good performance without losses.

Writing financial options is an explicit way to create an investment strategy that follows this pattern: by charging other for insurance against large market drops, you make modest returns most of the time, followed by sudden large disaster losses. Dynamic trading strategies can also show such option-like payoff patterns, and many other investment strategies share these characteristics, too. "Making markets" (that is, being the intermediary for buyers and sellers) seems to be one of them — dealers earn pennies on inventory for many years, until a sudden spike in volatility gets the dealer stuck with large losses.

Many funds don't even know that they follow strategies with such payoffs. The Great Recession of 2008-9 uncovered such patterns for many strategies previously believed not to suffer from them. Many investors (banks in particular) that had made small amounts of money for a long time suddenly lost it all. The momentum strategy in 2008 showed this pattern, too. Buffett had a funny quip here — "only when the tide goes out do you discover who has been swimming naked."

I don't think that academics or investment managers even understood the payoff patterns of their strategies, nor do I think that they tried to deceive their investors. They had just stumbled onto "it has made nice money for a long time" investment strategies. They followed them because they worked. My advice: be very skeptical about claims that someone expects to beat liquid financial markets.

The opposite of these "gathering pennies in front of a steamroller" strategies lose money most of the time but then gain a lot in a crisis. These strategies are very difficult to maintain. Which investors want to earn negative rates of return for years on end, while their peers are doing well? Only a few lucky bear investors can manage to maintain shorts. Tesla shares may eventually nosedive, but most Tesla short-sellers have already gone bankrupt.

So, yes, in theory, you can offer a fund with a negative market-beta strategy and low expected rates of return, because it provides great insurance that investors should want. In practice, your investors will drift away when the market goes up, and withdraw their gains when the market goes down and they have to cover their losses elsewhere. (I was on the board of a hedge fund that experienced exactly this.) It's tough to bet against the market, even successfully.

12.6 True Arbitrage vs. Risk(y) Arbitrage

Measuring investment performance brushes on a closely related topic — what exactly is the financial concept of arbitrage? Intuitively, an arbitrage is a great investment opportunity, perhaps so great that you should not be able to find one. Traders want to exploit any arbitrage opportunity that makes financial markets efficient as soon as it appears. It is a matter of basic financial literacy for you to understand what arbitrage is.

The Definition of Arbitrage

First, recall that the *law of one price* states that two identical items at the same time and location should have the same price. This is true in a perfect market, but even if the market is not perfect, it can be (and in fact usually is) still true. For example, even if all investors disagree about the future, even if there are taxes, even if there are transaction costs, and even if there is only one market maker, it should be, and usually still is, the case that one share of Intel Corp costs the same as another. But in a perfect market, the law of one price does not just *usually* hold; it must *always* hold. If it did not hold, you and the other infinitely many potential buyers could find arbitrage opportunities. The arbitrage concept is so important that you should understand it exactly, not just intuitively.

Begin Important

- A true arbitrage is a business transaction
  - that offers positive net cash inflows in at least some scenarios,
  - and under no circumstance — either today or in the future — has a negative net cash flow. This means that it is risk-free.

  An example: $5 for free.
- A risk(y) arbitrage is a business transaction that may not be risk-free but that still offers an excessive expected rate of return given its (risk and other) characteristics. A good way to think of a risk(y) arbitrage is as a great bet. Admittedly, the term "risk(y) arbitrage" is an oxymoron. However, Wall Street uses the term "risk arbitrage" for a particular type of trading (most often in the context of M&A transactions) that is similar to the sense in which we shall be using it. Thus, we shall commit the same sin.

  An example: A chance to win $1,000,000 with 99% probability and to lose $1 with 1% probability is not an arbitrage, though it would be an insanely great bet.

End Important

Arbitrage is an ex-ante concept, not an ex-post concept — beforehand, not after the fact. For example, it does not mean that a lottery ticket that won was an arbitrage. Ex-ante, a lottery ticket is not an arbitrage. Please also pay close attention to what the "no-negative-cash-flow" condition means in the definition of arbitrage:]subsect:FairBetsEx-ante fair bet

1. True arbitrage is not the same as "earning money without risk." After all, Treasuries do just that, and they are not arbitrage. The reason is that you have to lay out cash to buy Treasuries. This is a negative net cash flow today.
2. Arbitrage is also not the same as "receiving money today without a clear obligation to repay": If you are willing to accept risk, you can often receive cash today. For example, insurance companies take money in exchange for the possibility that they may have to pay up in the future.

Now contemplate the difference between the examples of the true arbitrage and the risk(y) arbitrage in the definition. You can lose $1 with 1% probability in the risky arbitrage, so it is "just" a great bet and not a true arbitrage. One difference is conceptual: Every investor would want to take a true arbitrage opportunity, but an infinitely risk-averse investor would not take a risk(y) arbitrage. This does not mean that, given an either-or choice, a less risk-averse investor would necessarily prefer the small, true arbitrage opportunity. In our example, would you prefer the $5 true arbitrage, if it cannot be repeated, to the risk(y) arbitrage with an expected payout of close to $1 million? (If you could scale the true arbitrage opportunity to take it infinitely many times, the true arbitrage opportunity would dominate.) Of course, this example of risk(y) arbitrage is extreme. More realistically, bets are never this great — "very good" is rare enough. And because there is still risk, you may not want to scale up good but risk(y) arbitrage bets in the same way you would always want to scale up true arbitrage bets as much as possible. Eventually, with enough investment in the risk(y) bet, your risk aversion would kick in and stop you from taking more of it.

Most of all, unless large public financial markets are very imperfect, you should expect not to find many great or arbitrage opportunities of either type there. If you agree with this assessment — basically that the world is sane enough not to have easy money grow on easily accessible trees where millions of others are passing by every day — then you can then draw some surprisingly strong conclusions about how these financial markets must work. If you disagree with this assessment, why are you still in this class? If you are right, you should be among the richest people in the world and there is little that this book and I can teach you.

However, if easy opportunities are scarce and money does not grow on trees, there are still plenty of important skills to learn. You can control risk; you can hedge and implement smart tax strategies; and you can find opportunities and superior projects in those markets that are not yet *too* perfect and *too* competitive. Look especially where you have a unique advantage, that will allow you to systematically outperform the others.

(Omitted solvenow)

More Hypothetical Arbitrage Examples

Of course, it is difficult to find real-world examples of arbitrage. Arbitrage is principally a conceptual issue. What would a hypothetical arbitrage opportunity look like? For example, if you can buy an item for $1, borrow at an interest rate of 9% (all costs, including your time), and sell the item tomorrow for $1.10 for sure, you earn 1 cent for certain today without any possible negative net cash flows in the future. If you ever stumble upon such an opportunity, please take it — it is a positive-NPV project! More than this, it is a true arbitrage because you cannot lose money in any scenario; it is riskless. Yet it is obviously not a very important arbitrage by itself. Searching for 1-cent arbitrage opportunities in financial markets is potentially more lucrative, because they often allow transactions to be scaled up. If you could repeat this 1-cent arbitrage 1 billion times, then you could earn $10 million. Unfortunately, although you may find an arbitrage that works once for 1 cent, it is unlikely that you can find such an arbitrage opportunity that works for 1 billion items. After all, you are not the only one searching in the financial markets! True arbitrage opportunities are difficult or outright impossible to find in the real world, especially in very competitive financial markets.

all arbitrages "matter"
nly if they are scalable.

Another hypothetical example of arbitrage would involve stock prices that are out of sync on different stock exchanges. If VFIAX shares are quoted for $51 on the Frankfurt Stock Exchange, and for $50 on the New York Stock Exchange, you could theoretically buy one share in New York for $50 and sell it in Frankfurt for $51. You then pocket $1 today. If you can do this with 20,000 VFIAX shares worth $1 million, you could earn $20,000 without effort or risk.

But before you conclude that this is an arbitrage, you still have to make sure that you have not forgotten about costs or risks. The arbitrage may be a lot more limited than it seems — or may not even be present at all. Consider the following issues:

1. Could the price change in between the time you buy the shares in New York and the time you sell the shares in Frankfurt (even if it is only 3 seconds)? If such execution-timing risk exists, this is not pure arbitrage because there is a chance of a negative net cash flow. The real-world evidence suggests that price discrepancies between markets often disappear within a few seconds.

2. Did you account for the direct and indirect transaction costs? How much commission do you have to pay? Is $51 the Frankfurt bid price at which you can sell shares in a market, and $50 the NYSE ask price at which you can buy shares? Can you sell the share in Frankfurt and get it quickly enough from New York to Frankfurt to make the closing? Have you accounted for the value of your own time watching the screen for opportunities?
3. Could the share prices move when you want to transact a significant amount of shares? Only the first 100 shares may be available for $50 for a net profit of $100. The next 900 shares may cost $50.50 — perhaps still worthwhile, but less profitable. And buying the remaining 19,000 shares may cost you $51 or more.
4. Did you account for your fixed cost of setting up your business? If it costs you a million dollars to get offices and computers in order to "arbitrage" a few thousand dollars, it is obviously not a real arbitrage. So you must account for how expensive it is to set up your operations.

It may be that small arbitrage opportunities occur from time to time, but large financial firms are constantly running automated computer trading programs that search for even tiny arbitrage opportunities in order to exploit them as soon as they appear — and thereby make them disappear.

(Omitted solvenow)

Limits to Arbitrage

(Omitted anecdote)

There is yet another problem with near-arbitrage. Let's say that you have found a discrepancy between the value of two assets and you are sure they will converge *eventually*. You would need to buy the cheaper version and go short on the more expensive version. There are, of course, the transaction costs of doing so, especially shorting. But let's say that they are small.

The bigger problem is different: If the relative valuation makes no sense at 10 bps, what stops it from going to 20 bps, at least for a while? If you have to borrow money in order to make this bet, the temporary widening of the near-arbitrage could even kill you.

Let's go back to Bitcoin. As I already wrote, I believe that Bitcoin is a bubble. I see no value at the end of the chain. I am positive that Bitcoin will not become a major world currency. When speculators collectively realize that there will be no future suckers extrapolating the past and hoping for free money, the speculators will disappear, too. Bitcoin will collapse. However, given that Bitcoin now (at the end of March 2022) trades for $40,000, what prevents it from trading for $400,000? One makes (almost) as little sense as the other.2]subsect:bitcoinBubbles and Bitcoin

There is, of course, also the other "little" problem: it is not possible to reliably short Bitcoin. Bitcoin trades mostly on exchanges that will themselves go bust if Bitcoin suddenly collapsed. This means that they are not likely to be able to pay up if the short bet pays off a lot.

I just hope that by the time I write the next edition of my book, I will not have to double up on my prediction of Bitcoin's demise, with Bitcoin trading for $400,000. I am never wrong, just like Larry David's Super Bowl commercial for the FTX crypto exchange.

12.7 Corporate Consequences of Market Efficiency

How does the EMH matter to you if you are a manager? Does it matter whether financial markets are perfect, efficient, or neither? Because a perfect market implies an efficient market, you need to think about three different cases:

1. The market is efficient and perfect.
2. The market is efficient but not perfect.
3. The market is neither efficient nor perfect.

These cases help you organize your thinking about what it takes to create value — which is *the* most important question if you are the CFO. Can you add value by changing your capital structure? Can you create value by splitting your shares, so that every share becomes two shares? Can you create value by paying out dividends next year rather than this year? Can you create value by changing how you present your earnings to investors? Can you create value by taking over other companies when they are priced too low if you do not have any unique knowledge or anything unique to add?

If the Financial Market is (Close to) Perfect

If the financial market is perfect, the answers to these questions are simple — they are always no. It does not matter how the firm communicates its earnings to investors, what its capital structure is, how many shares it has, how it pays out its dividends, and so on. In fact, you already know that the firm is worth the value of its underlying projects' present values. Everything else is irrelevant.

Earnings reporting:]sect:financialstmtsDo reported earnings matter? For example, if you have previously reported your foreign division's earnings separately and now you consolidate them into your main earnings, you will indeed increase the firm's reported earnings. However, it will not create anything intrinsically valuable. Such a change should not increase or diminish firm value. Your firm owned the subsidiaries' cash flows before and after its reporting change. Your investors can add or subtract the subsidiaries' numbers themselves, whether you include or exclude them in your overall report.

Capital structure: .75]sect:mandmformalCapital structure arbitrage For example, say your firm is currently financed with equity only and worth $100, but if you had a 50-50 debt-equity ratio it would be worth $102. In this case, an arbitrageur could buy your firm, issue $51 in debt and $51 in equity, and pocket $2. With legions of arbitrageurs competing to do this, your firm value would instantly adjust to $102. Thus, a $100 price for your firm would be absurd.

Stock splits: .5]sect:perfectmktirrelStock splits In a stock split, each old share becomes multiple new shares. For example, if each share trading at $80 were to become two shares, the new shares should trade for $40 each in a perfect market. Nothing fundamental about your underlying projects would have changed. Splitting by itself cannot add value. If this were not the case — for example, if shares would be worth $41 each after the split — arbitrageurs would buy the old shares for $80, and sell them an instant later for the equivalent of $2 \cdot \$41 = \$82$, pocketing $2.

Dividends:]sect:perfectmktirrelStock dividends The same argument applies to dividends. In a perfect market, a $100 firm that pays $10 in dividends should be worth $90 thereafter — no value is magically created or destroyed. Keeping the money for another year in the marginal zero-NPV investment (e.g., Treasuries) is as good as paying it out. Investors in a perfect market can borrow against this extra future money and use it today.

The lesson is simple: As a manager, you should forget about the smoke and mirrors and instead focus exclusively on finding and executing projects with positive net present values.

If the Financial Market is Not Perfect but At Least Efficient

If markets are not perfect but efficient, the implications are not as profound. However, it means that you can still obtain valuable market intelligence. Your market price is the aggregate assessment of many investors who have put their money where their mouths are. The market price aggregates a whole lot of information that you as a corporate manager may not learn as easily yourself. For instance, if your stock price seems very high relative to current fundamentals, it probably means that the market sees great opportunities ahead for your firm and expects that you will take them. Thus, you should consider growing the business. Naturally, a high firm value allows you to raise more funds from the financial markets at favorable rates. On the other hand, if the stock price is very low, it probably means that the financial market anticipates your business to go down or expects you to waste the remaining money. In this case, you should think carefully about whether you should reinvest investors' money into the business or into repurchasing the (relatively cheap) stock.

In addition to learning from your own company's market price, you can also learn from all sorts of other market prices. You can find out how good your competitors' opportunities are, and whether you should get into the fray. Commodity price information can also be very helpful. If the price of oil in the forward market is $100/barrel, it probably does not make sense for you to plan ahead based on an oil price of $70/ barrel. The financial market price for oil forwards is very large and efficient. It makes no sense for you to plan your business around much lower or higher oil prices in 6 months, simply because if you really knew this better, you could get rich easily without needing any of your current businesses — just start trading oil futures. This may sound obvious, but it is sometimes easy to overlook the obvious in the heat of battle. For instance, a friend of mine who sat on the corporate board of a large conglomerate oil company objected to its capital budget. The company planned to explore for more oil, based on a working assumption that oil prices would recover and double within two years. Whether the projections would turn out to be true or false was actually irrelevant. This oil company could just have purchased oil in the market much more cheaply instead of drilling for it. Why bother exploring for oil if you can buy a confirmed reservoir as cheaply in the market? Eventually, they came to their sense. If you are a farmer planting, the futures exchanges provide you with forward prices for corn and wheat, and you can use this free price information to help you decide which crop to plant.

Let's consider a specific example of how you can learn from market prices in an efficient market. Put yourselves in the shoes of a smart and successful manager of an aircraft manufacturer. Every morning, you read the newspaper, and every morning you think that company X should really be worth a lot more. It makes no sense to you that X has annual earnings of $10/share but its shares are trading at only $50/share. X just seems undervalued. Should

you go out and buy it? If the market is perfect, the answer is no. You would have no competitive advantage in owning X. Hordes of professional arbitrageurs would have traded and eliminated the opportunity in an instant, and less expensively than you ever could. On the other hand, owning X would not do any harm, either. But let's take away the perfect market assumption and leave only the efficient market one. This means that both your aircraft company's price and the price of X are correct. Buying X because you think that X is undervalued is likely to be wrong. After all, our working assumption is that the financial markets have used all available information to find the best possible price.

However, in the absence of perfect markets, the efficient market does not mean that you should never be able to create value by buying other companies. You can indeed sometimes create value. The trick is that you must be able to do something that investors cannot do for themselves, because the market is imperfect. Most likely, this would be related to your business's real operations. For example, if X is a supersonic aircraft parts supplier, you may have better information about the supplier's product. Unlike anyone else, only you may know that you will reward it with a huge contract soon. Or, by owning the patents of this supplier, you may make it more difficult for other aircraft companies to compete with you. Or you may find cost savings by cutting out the middleman in purchasing these parts, or by improving X's products through your own intellectual capital, or by increasing the scale of operations. All of these strategies can add value to the firm — value that outside arbitrageurs cannot accomplish without you. (This kind of unique ability to create value violates the "perfectly competitive market with infinitely many potential buyers" assumption of a perfect market.)

But be careful: Market efficiency means that you cannot create value for your shareholders simply by your personal view (without special information) that X is undervalued. Yes, you may be smart, but the financial markets are just as smart and presumably could recognize just as well whether X is undervalued — in fact, chances are that the target was rightly valued to begin with and it was you who got the target value wrong. For example, if you manage a company manufacturing aircraft parts and X manufactures pharmaceuticals, it is highly unlikely that you would create value for your shareholders by buying X, even if the firm X is trading for only 5 times earnings and this seemingly-low valuation makes no sense to you.

The same argument applies to all sorts of other corporate actions. You may be able to create value by reducing perfect market barriers. For example, you may be able to create value by reducing the costs that investors incur when trading your shares (e.g., by listing on an exchange). Or you may be able to reduce the mistrust that your investors might have in your creditworthiness by hiring a good auditor or by reporting your earnings in a transparent fashion. Indeed, there is evidence that many corporate activities can create value by reducing the perfect market frictions, even in very efficient financial markets. For example, when firms split their shares 2-to-1, it is not necessarily the case that the two post-split shares are worth exactly half of the pre-split share of, say, $80. Instead, they tend to be worth a little more, say, around $40.20. The likely reason is that managers signal their confidence in the future by splitting shares today. This sort of corporate move brings more information to the market. Importantly, it is not the split per se that makes the share price go up, but the expected increase in future real cash flows that does.

If the Financial Market is Not Even Efficient

Loosely speaking, financial markets tend to be reasonably, but not always perfectly, efficient. Perfect market efficiency is almost surely *not* a good description of reality. Even in a perfectly rational market, as an executive, you may know the firm value better than the market — for example, you may know that your company is likely to sign a large contract, but this information cannot yet be disclosed. What should you do if you know that the stock price is not equal to the appropriate market value? The right way to conceptualize your problem is to consider what you would do if you were the sole owner of the firm. You would really care about firm value. (As its executive, you should want to maximize this value on behalf of the owners.)

If your shares are undervalued, you should recognize that your cost of capital is effectively too high, given the true characteristics of your project. The reason is that you cannot raise risky capital at fair prices — especially equity capital. The CAPM clearly is no longer the right model for the cost of capital.

For example, assume you know that you are a startup whose current projects will return $500 tomorrow but you are not yet able or legally allowed to disclose this. The problem is that without this information, outside investors may value your firm only at $150. Also assume that you have no cash and that you can only raise financing through selling more equity to potential outside investors. Now assume you come across a new project that costs $100 and will return a terrific $200 tomorrow. The problem is that your investors do not think that the firm will return $700, falsely believing that the combined firm will only be worth, say, $250. Thus, to raise the additional $100 in capital that you would need, you would have to sell 40% of your firm (truly worth $280), and keep only 60% (truly worth $480) of the true $700 return. $You would therefore be better off passing2]subsect:ambiguousvalueinimperfectmarketsSeparation of financing and investing up this new project and just taking the $500 from the old project. Put differently, your true opportunity cost of new capital (selling something worth $280 for $100) to fund this project is way too high for you.

You would definitely not want to raise cash at these "high" prices. Instead, you would want to do the opposite. The best use of corporate cash may now be to repurchase your own cheap, underpriced shares — for example, from other investors. However, there is an intrinsic paradox here: As an executive, you are supposed to act on behalf of your shareholders. Therefore, repurchasing underpriced shares from them at bargain prices would not be what would make the selling shareholders better off. (It would, however, make your remaining shareholders better off.)

If your shares are overvalued, your cost of capital would be too low. You should be tempted to take more projects. This is easiest to see if you again consider what you would do if you were the primary owner of this overpriced firm. You would want to sell more equity shares at higher prices and pay the money out in dividends to existing shareholders. (Alternatively, you can just invest in Treasury securities.) Here the paradox is, of course, that just one instant later, as CEO, you are now the representative of these new shareholders to whom you have just sold overpriced shares. They will not be happy campers. (Many researchers believe that this is exactly what happened when AOL purchased Time-Warner at the height of the Internet craze in the late 1990s. AOL used its overpriced shares to buy Time-Warner's real assets.)

These are robust insights for corporate managers who are not conflicted and wish to act on behalf of their existing shareholders. **Begin Important**

When managers have superior information:

- If the firm is undervalued, CEOs should assume a relatively high cost of capital and consider repurchasing the firm's own shares.
- If the firm is overvalued, CEOs should assume a relatively low cost of capital and consider issuing more of the firm's own shares.

A good decision rule for managers is to take projects up to the point where the marginal costs and benefits of projects are the same as what they could obtain from repurchasing or issuing the firm's own shares. **End Important**

(It can become a bit more complex if you see yourself as a representative of both new and old shareholders, though.)1]sect:humanbiasOverconfidence But be careful: Most executives are notorious for *always believing* that the financial markets do not fully reflect the value of their companies even if they have no inside information — as an executive, you should be wary of your own perceptions and biases!

(Omitted solvenow)

Comparison of Market Concepts

Here is a summary of the two conceptual classifications of how markets work:

Efficient versus inefficient markets: If the market is efficient, you can learn from financial market prices, because they accurately incorporate the information of financial market participants. This means that you cannot create value by buying other companies just because you think that these companies are worth more than they are trading for.

If the market is inefficient, you may be able to identify underpriced firms that you can take over, and/or create value by improving how and when information reaches the market.

Perfect versus imperfect markets: If the market is perfect, you can focus exclusively on your projects' net present values. You can forget about most financial choices, such as what your capital structure should be, how you should report earnings, and so on.

If the market is imperfect, you can create value, often by reducing the market imperfections themselves. For example, you could signal what you know about your company's prospects by reporting earnings sooner. On occasion, this can even become a dilemma: For example, what should you do if you know that a project has a positive NPV but the financial market does not believe you? If you take it, your stock price may go down. Now you have to think about the lesser of two evils — passing up on the good project vs. passing up on a higher stock price.

In the real world, financial markets are definitely not 100% perfect. For large firms, they are very close to efficient, but this is not necessarily so for small firms. Still, the economic magnitudes of deviations should be fairly modest. As a real-world manager of a publicly traded corporation, you are generally better off focusing on underlying value creation than on actions that investors can accomplish for themselves without you. It makes sense for you to believe that market prices are almost always informative, but not to believe too slavishly that they are also always fully efficient — you may have better information than the market. Use it wisely when you have it.

12.8 Event Studies

The immediacy of price reactions in any efficient market offers a surprisingly useful real-world application: In some cases, market price reactions can allow you to estimate value consequences more easily than traditional NPV techniques, through the use of a technique called an event study. An event study is an empirical analysis of the effect of a set of events on the prices of assets. The idea of an event study is that if the public market is valuing projects appropriately, and if the value of an unexpected event or action is $1 million, then the stock price should increase by $1 million at the instant the event becomes publicly known. You can therefore (often) back out cash flow value changes from stock price changes. The details of how to conduct such a study are in the appendix.

Capital-Structure-Related and Other Event Study Results

Researchers have run event studies on all sorts of interesting events, ranging from new legislation, to corporate name changes, to analysts' opinions, to corporate earnings, to stock splits, to corporate dividends, to corporate debt and equity issuance and retirement, to deaths of founders, and so on. Here are some of the more important findings. (You will see some more evidence obtained from event studies again in later chapters, especially in the chapters on capital structure and payout policies.) On the day of the announcement, firm values *increase* on average:

- When firms announce increases in dividends, share repurchases, or stock splits (by about 0.1-1%; if you are interested, there is a longer explanation in Chapter).
- When firms are taken over by other firms (by about 10-30%).
- When firms announce earnings that significantly beat analysts' expectations.
- When pharmaceutical firms announce that the FDA has approved one of their drugs.
- When the founding CEO dies (by about 3-4%).

Conversely, firm values *decrease* on average:

- When firms announce new stock sales (by about 1-3%).
- When firms overpay for other firms in acquisitions.
- When firms announce lower-than-expected earnings.
- When firms fend off an acquirer who has made a bid.
- When drug firms announce that the FDA has rejected one of their drugs.

In both cases, because researchers usually do not know the markets' probability assessments prior to these announcements (some of the effects would have already been anticipated and thus already incorporated in the stock price), the true all-inclusive value changes of these events are likely to be more extreme than the observed event-study responses. This makes the empirical effects conservative lower bounds.

Event studies have also informed us whether certain government regulations had a positive or negative impact on firms. For example, we know which firms were helped and which were hurt when the telecommunications, trucking, and airline markets were deregulated — or how the Treasury's rescue program of 2008 ("TARP") helped some banks, but not others.

(Omitted anecdote)

(Omitted solvenow)

Summary

This chapter covered the following major points:

- Market efficiency means that the market uses all available information in setting prices to offer "appropriate rates of return."
- In the short run, the appropriate expected rate of return on stocks must be small. Therefore, market efficiency prescribes that stocks roughly follow random walks.
- In the long run, it is rarely clear what this "appropriate rate of return" should be. Because noise makes it difficult to measure the average rate of return accurately, it is also difficult to test either models like the CAPM or long-run market efficiency.
- Beliefs in efficient markets come in different forms.
  - The standard efficient markets classification emphasizes what information it would take to beat the market: weak form (past stock price patterns are not enough to beat the market), semistrong form (other historical firm information is not enough to beat the market), and strong form (inside information is not enough to beat the market).
  - A more current efficient markets classification emphasizes the rationality of the stock market: true believer (stock prices always reflect underlying project NPVs); firm believer (small deviations between price and value, but difficult to take advantage of); mild believer (small deviations between price and value, and somewhat possible to take advantage of); or nonbeliever (arbitrage opportunities abound).
- The overall evidence suggests that it is not easy to become rich by exploiting financial market inefficiencies — a belief shared by most finance professors. The relative strength of their beliefs in market efficiency — the extent to which professors believe that market prices always reflect underlying value — separates finance professors into "rationalists" (or "classical" economists) and "behavioralists."

😄 Touché: Something I overheard at a conference: A famous hedge fund manager asked "if you are so smart, how come you are not rich?" — to which the equally famous finance professor responded "if you are so rich, how come you are not smart?"

- In a perfect and efficient market, investors should not find arbitrage opportunities:
  - True arbitrage is a riskless bet with no negative net cash flows under any circumstances. Everyone would like to take all true arbitrage opportunities. When and if they appear, they are likely to be very small and to disappear quickly.
  - Risk(y) arbitrage is more like a great bet. An infinitely risk-averse investor would not want to take it, because there is a chance that risk(y) arbitrage will lose money.
  - Both true and risk(y) arbitrage opportunities should be very rare in the real world. An investor who is not too risk-averse may or may not prefer taking one large, great bet to taking one tiny, true arbitrage.
- Given the millions of investors, many will beat the stock market by chance, and some investors will beat the stock market many years in a row. Market efficiency does not mean that there are not some investors who will beat the stock market 10 years in a row *ex post*; rather, it means that any one particular investor is unlikely to beat the stock market *ex ante* 10 years in a row.
- Managers can learn valuable information from market prices, both from their own share prices and from other prices. To improve corporate firm value, managers must create fundamental value — they must undertake positive-NPV projects. Simple uninformed activities such as purchasing another firm to lower risk or splitting shares will not add a lot of value — if any.
- Event studies allow you to ascertain the corporate value impact of sudden events, such as election results, regulatory action (FDA rulings), or corporate events (dividend increases).
