

A First Look at Investments

Background

The subject of investments is so interesting that I first want to give you a quick tour, instead of laying all the foundation first and showing you the evidence later. I will give you a glimpse into the world of historical returns on the three main asset classes of stocks, bonds, and “cash,” so that you can visualize the most important patterns — risk, reward, and covariation. This chapter also briefly explains how to empirically distinguish causality from correlation, and it describes a number of important institutions that allow investors to trade equities.

7.1 Stocks and the S&P 500 Index

Investments are often classified into broader asset classes, like cash, bonds, stocks (also called equities) — or precious metals and art. To summarize their performance, there are indexes that describe their value evolution. These indexes are often based on specific portfolios and give investors basic guidance as to their historical performance. Let's start by describing U.S. stocks.

The most popular U.S. stock-market index is the S&P 500 ([S&P 500](#)). (S&P stands for Standard and Poor's, the firm that invented this index in 1923 and continues to maintain it.) Although not precisely accurate, you can think of its constituents as the 500 largest publicly-traded firms (out of about 4,400), all weighted by market-value (value-weighted) in one portfolio. At the end of 2021, the [S&P 500](#) index represented about 85% of the market capitalization (i.e., the total value — the number of all outstanding shares times the price per share) of all publicly traded stocks. The largest firm, Apple, had a market capitalization (commonly referred to as market cap) of just under \$3,000 billion dollars at the end of 2021. The smallest firms had market caps of about \$70 billion.

The [S&P 500](#) index is so prominent that nowadays about 20% of all capital in the stock market is invested through funds that mimic the [S&P 500](#) portfolio — and many more investment funds are performance-benchmarked to it. Offering investors an easy way to buy, sell, or simply speculate on the [S&P 500](#) is one of the most competitive businesses in the world. (Some vendors, like Vanguard, even compete with themselves, offering multiple funds that all try to mimic the [S&P 500](#).)

However, there is one complication to using the index as a stand-in for the stock market: it is a value-weighted (also referred to as a capitalization-weighted) price index, which explicitly ignores dividends. Thus, the rate of return

on *any* fund holding the same 500 stocks (in the same proportion as the index) is higher than (the percent capital gain change in) the index itself. The actual holding fund earns the dividends, too. As a result, when real-world funds use the [S&P 500](#) as a benchmark, it makes their performances look relatively better. (Nevertheless, the evidence suggests that less than half of all funds have actually beaten even this low benchmark — 1]ch:effbehavBeating the Marketit's hard to beat the market.)

You can easily notice the difference between the annualized percent gain on the [S&P 500](#) index and the annualized rates of returns on a few prominent funds that mimic the [S&P 500](#) index. For example, over the last 10 years, the annual mean rates of return were

	Quoted	Actual Funds			
	Index	.. Vanguard ..		Fidelity	SPDR
	S&P 500	VOO	VFIAX	FXAIX	SPY
2012 - 2021	14.09%	16.31%	16.35%	16.38%	16.28%

As just mentioned, the stark 2.2% return difference was mostly due to the dividend yield. The differences among the four actual index funds were minute. Moreover, they all had greater than 99.8% correlation on daily returns with one another (and with the [S&P 500](#) index). For purposes of assessing performance, any of the four would do. We shall use the Vanguard [VFIAX](#) fund in what follows, because it has been around the longest.

Graphing Historical Returns

(Omitted fig)

Let's start with Figure ?? . It shows the year-by-year rates of return of [VFIAX](#) . The plot and the table illustrate the same data: For example, you would have lost 37% in 2008, gained 15% in 2010, and so on. The arithmetic average rate of return over the entire 17 years from 2005 to 2021 was 12% per annum.

These particular years were marked by two unusual aspects. First, there was one very steep downturn in 2008 (the start of the Great Recession). Second, there were unusually few negative market returns. At least since the late 19th century, stock-market investors typically suffered negative returns in about 1-in-4 years (and rates of returns below the risk-free rate in about 1-in-3 years). The lower frequency of modest negative returns means that, despite the market crash in 2008, the last 17 years were exceptionally good for stock investors.

(Omitted fig)

Figure ?? takes the same data but presents it differently. It shows a type of histogram that is based on the number of returns that fall within a range. This plot makes it easier to see how spread out returns were — how common it was for the [S&P 500](#) to perform really badly, perform just about okay, or perform really well. The histogram shows that the most common rate of return was between 15% and 25% per year — very good. The standard deviation,3.5]subject:uncertainty-variance-firsttimeUncertainty and Variance the most common measure of risk, is related to how wide the histogram is. Here, it was 24% per annum. Unusually, over these last 17 years, the rates of return did not appear normally distributed. They were more bimodal. Over longer periods, the shape looks more like a bell curve. 1]app:bellmo150 Years of Monthly Stock Returns

Compounding: Arithmetic and Geometric Rates of Return

(Omitted fig)

Naturally, most investors are more interested in how much money they made than in statistics.

With the mean and standard deviation of these 17 annual returns, can you now determine how much money you would have earned for a \$1 investment? Specifically, could you take \$1 and the 12% arithmetic average return and use the compounding formula? Compounding arithmetic average returns would indicate a final wealth of $\$1 \cdot (1 + 12\%)^{17} \approx \5.87 . Unfortunately, this would also be wrong (and it is an error all too commonly committed by investors).

Instead, you need to compound the individual rate of return. The result is shown (on a logarithmic scale) in Figure ?? . For example, over the first four years, the compound return of \$1 invested on January 1, 2005, would have been \$0.81 on December 31, 2008 — a loss of 19%:

(Omitted eq)

However, the stock market did well after 2008. By the end of the sample, the \$1 investment starting in January 2005 would have turned into \$5.53 in December 2021.

Why was \$5.87 incorrect? Think of an example. If you had earned a rate of return of -50% (you lose half) followed by $+100\%$ (you double), your compounded rate of return would have been zero. However, your average rate of these two returns would have been a positive $(-50\% + 100\%)/2 = +25\%$. Equivalently, if you had earned $+50\%$ followed by -50% , you would have ended up with only $1.5 \cdot 0.5 = 75\%$ of your investment, a negative rate of return (-25%). (There are many real-world examples in which the compound rate of return was -100% [you lost all your money], yet the average rate of return was still positive. Yikes 😱!) And if the difference between \$5.87 and \$5.53 seems small to you, realize that [VFIA](#) (S&P 500) alone has over \$250 billion under management. A mere 6% thereof comes to \$15 billion! Financing requires attention to detail, because many financial markets are so close to perfect that even small opportunities can often scale up tremendously.

The *annualized* compound rate of return is called the geometric average. To compute it, you must annualize the total investment return the way you learned it in Chapter . With an investment of \$1 and a final value of \$5.53, the annualized rate of return was

(Omitted eq)

This 10.6% geometric average rate of return is less than the arithmetic average rate of return of 12.0%. The way to interpret this 1.4% discrepancy is as follows: A risk-free asset without volatility and a 10.4% rate of return each and every year would have offered the same final return of \$5.53.

The wedge between the arithmetic and geometric rate of return is caused by the volatility of the returns. Comparing assets with different volatilities on the basis of their historical *arithmetic* average rates of return can lead to misleading conclusions. Riskier assets can have higher arithmetic average rates of return and still underperform

*Incidentally, all our annual returns were themselves compounded from daily returns.

safer assets (and this is not on a risk-adjusted but on a plain basis). Be careful — funds usually quote arithmetic average rates of return. This makes riskier funds look relatively better (e.g., when comparing stock funds to bond funds).

Unfortunately, this is not just one of those “academic egghead” concern. Many, if not most, investors mistakenly assume that the historical arithmetic rate of return is a fair representation of what they can expect to earn every year in the future. But they will not earn the arithmetic rate of return. Instead, to the extent that historical returns are indicative of future returns (and this is not necessarily true, either), investors should expect both the historical arithmetic rates of return and the historical variances to continue and therefore expect not average but geometric rates of return.

There are many other tricky conceptual aspects here. For example, would you rather invest your money (a) into a gamble in which you make +60% and –50% with equal probability at the same time, or (b) into a gamble in which the returns come sequentially, making 60% half the time and losing 50% the other half the time? In the former case, you expect to end up with

(Omitted eq)

In the latter case, you expect to end up with

(Omitted eq)

Given a choice, go for the simultaneous investment over the sequential investment.

This problem can also bite in NPV applications, where managers have to estimate equivalent rates of return both for their projects and for equivalent projects. Managers more commonly and incorrectly rely on (historical) arithmetic average rates of returns. This may not be terrible as long as their projects are equally risky as their benchmarks — but this is not a given. You have been warned.

As already explained, the geometric average rate of return is always less than the arithmetic average rate of return, except when there is zero time volatility. The more risk, the bigger the difference. You should compare assets with different volatilities only based on their geometric rates of return. Nevertheless, a “lazy rule of thumb” approximation formula often works ok: If rates of return are approximately normally distributed (like a bell-shape) with a modest variance only, then the arithmetic mean is higher than the geometric mean by about half the variance. In our example, the [S&P 500](#) had an annual standard deviation of 24%, which comes to a variance of $24\%^2 \approx 576\% \approx 0.0576$. Thus, the approximation formula says that the geometric rate of return should have been about 2.9% lower than the arithmetic return. In these 17 years, the approximation overshot — the actual difference was only 1.4%. The problem here is that our stock return series did not look normally distributed at all — it had one really bad outlier, 2008, which jacked up the volatility.

There is one further novel aspect to Figure ??, which is the gray-shaded area. It marks the cumulative CPI inflation..5]sect:apples-nominal-inflationApples, Oranges, and Inflation The purchasing power of \$1 in 2005 was about the same as \$1.39 at the end of 2021. Thus, the \$5.53 nominal value in 2021 was really only worth $\$5.53/\$1.39 \approx \$3.98$ in 2005-inflation-adjusted dollars. (And,]sect:taxbasicsTax Basics of course, none of these figures take income taxes into account.) Still, this was an excellent real rate of return by historical standards!

(Omitted solvenow)

7.2 Stocks, Bonds, and Cash

Let us expand our perspective to fixed-income investments, which promise to pay out a set level of cash flows to investors. In this section, we first look at two of the most common types — “cash” and “bonds” — over the same 17 years as we did in the previous section (before we return to stocks).

The designation *cash* is a slight misnomer, because it does not designate physical dollar bills. Instead, it still means debt securities — financial instruments that require the borrower to pay back the face value of the loan plus interest income — but very short-term, liquid, and low-risk. Money market fund[Money market funds] are prominent cash-type investments; they are loans that governments and corporations sell to investors.[†] Unlike the short-term money market funds, bonds are debt instruments that have longer maturity than cash. You already know a lot about bonds and their many different varieties from Chapter . When the term is not further qualified, most investors mean long-term U.S. Treasury bonds.

Just as we used the Vanguard **VFIAX** mutual fund as a stand-in for the 500 largest stocks, we can use the Vanguard **VMRXX** money-market fund as a stand-in for cash and the Vanguard **VUSUX** long-term Treasury bond fund as a stand-in for bonds.

Of course, you should remain aware of the limits of broad asset-class categorizations. They are not representative of *each and every* asset that would seem to fit their designations. For example, some long-term corporate bonds by financially distressed issuers behave more like equities. Similarly, some firms may own a lot of bonds, and their equity rates of return could look more like those on bonds and not like those on stocks. (In fact, a Vanguard bond fund is like a company that holds bonds. By buying “stock” in the Vanguard bond company, we obtain bond-like rates of return.) It would also be perfectly reasonable to broaden or narrow the types of investments in these three asset classes. (We would hope that such modifications would alter our insights only a little bit.)

(Omitted fig)

Let’s compare the rates of return patterns on stocks, bonds, and cash over our 17 years. Figure ?? plots exactly the same types of graphs as those in Figures ??, ??, and ??.

The top row shows how tight the distribution of cash returns was around its average rate of return of 1.4% per year. The standard deviation was only 1.8% per year. You would never have lost money in nominal terms, but you would rarely have earned much, either. Your investment portfolio value would have slowly and steadily marched upward. Each dollar invested on January 1, 2005 would have become \$1.27 at the end of 2021.

Of course, inflation would have eroded the value of each dollar. In purchasing power, you would have lost money. (It would have been worse for many retail investors, who would have had to pay personal income tax on the interest of \$0.27. If their tax rate was 40%, they would have been left with \$1.16 in nominal terms, equivalent to $\$1.16/1.39 \approx \0.83 in inflation-adjusted terms.)

[†]Other investments that can be used as cash are short-term certificate of deposits (CDs),¹subject:how-banks-quoteCDs savings deposits, or commercial paper. (These are also briefly explained in the Book Appendix.)

The middle row shows that long-term Treasury bonds were riskier than cash. The histogram is much wider and the standard deviation is 12.2% per year. Fortunately, in exchange for carrying more risk, you would have also enjoyed the higher average rate of return of 6.9% per year. And \$1 invested in 2005 would have become not just \$1.27 but \$2.82 — again before taxes. (If you had to pay 40% taxes on \$1.82, you would have been left with about \$1.10 in interest. The \$2.10 in after-tax return would have been worth about $\$2.10/1.39 \approx \1.51 in inflation-adjusted 2005 dollars.)

The bottom row repeats the information on the S&P 500 for comparison. Stocks were riskier than bonds. The stock histogram was more “spread out” than the bond histogram, primarily due to its large outlier in 2008. However, stocks also had many more years that were great. Thus, the higher risk of stocks would have paid off. As we already discussed, the standard deviation of 16.7% per year was compensated with an average mean rate of return of 12.0% per year. Your \$1 invested in 2005 would have ended up being worth \$5.53 in 2021 — again before taxes. (Stocks are also taxed less aggressively than fixed-income securities. We will return to this subject in Chapter .)

7.3 Many More Asset Classes

(Omitted tbl)

(Omitted tbl)

There are also many other asset classes available to investors. Even our three main asset-classes are often sub-classed into smaller categories. Table ?? lists a few of them. In our 17 years, assets that were riskier also tended to deliver higher average returns — though not all did.

Among different types of firm-size-based stock-market index portfolios, there was not much difference in how firms performed. Small-cap, mid-cap, and large-cap stocks all earned about 10-11% (geometric) — with differences that were not statistically meaningful. (Reminder: Cap is a common abbreviation for “market capitalization,” itself a fancy way of saying “market value.”)

“Value” stocks (usually older firms with high book values relative to market values) underperformed “growth” stocks (usually younger firms with the opposite pattern) by about 3% per year. This outcome is modestly ironic, because until the mid 1990s, value stocks had outperformed growth stocks — a pattern that was often attributed to the brilliance of value investing, an approach pioneered by Graham and Dodd in the 1920s, and thereafter picked up and expanded upon by Warren Buffett, famous finance professors Fama and French, and countless quantitative hedge funds. The “value effect” was also a flagship pattern in behavioral finance — the supposed factoid being that naive investors incorrectly flock more to flashy growth stocks and ignore mundane value stocks, thereby allowing the neglected latter to outperform the popular former. Well, it’s always dangerous to extrapolate in financial markets! Investors may not have remained as ignorant when the evidence punched them in the face. By buying a lot more value stocks, they would have driven up price and lowered value stocks’ expected rate of return — and thereby eliminated the value anomaly itself.

The next set in the table are funds that invested into industry categories. REITs are themselves real-estate investment trusts, i.e., investment vehicles that purchase and lease out real-estate. The table shows that Vanguard’s fund of many REITs ([VGSLX](#)) performed a little worse than the stock market overall, but not badly so. Real-estate

was a good investment during these 17 years. Not all industry portfolios did so well, though — for example, Energy (VGELX) and Financials (VFAIX). To be clear, at the outset in 2005, no one suspected that these sectors would perform so much worse than other sectors (and in particular tech stocks). Do not necessarily expect this pattern to repeat, either. We will return to the concerns about extrapolating historical returns later in this chapter.

The next rows show that the U.S. stock market did much better than its European and Pacific cousins. The latter delivered only about half as much return as U.S. stocks. Again, there is no law of nature that says that this pattern must continue — or even that stock markets will have a positive compound rate of return over the next 17 years. Then, again, they could do even better.

Table ?? expands our universe. For bonds, the patterns were generally more predictable. Typically, the riskier the bonds (e.g., junk bonds and long-term bonds), the more yield they promised and ultimately delivered.

There are also many other financial and non-financial investment classes. For example, precious metals and in particular gold have been investments for millennia. Fine wines have been another investment, although only for a few hundred years — and perhaps utilized more commonly among connoisseurs with a reputation for pretense. Over these 17 years, gold and wine performed similar to stocks, earning about 8.4% per year on average with a risk of 15-17% per year.

However, investments in commodities (energy, grain, metals, livestock, etc.) performed poorly — *losing* money even in nominal terms. Investments in *Fine Art* barely broke even. And be suspicious — the art and wine indexes were not actually traded but just indexes. And they are published by organizations with interests in promoting their uptakes as investments. The true returns may have been even lower.

When you looked at the table, did you look at the arithmetic or the geometric average rate of return? The first columns in tables are always catchier than middle columns! I warned you — it is easy to forget that it is the latter geometric returns that matter, not the arithmetic returns! For example, energy stocks did *not* outperform Treasury bonds, even though they offered higher average rates of return.

(Omitted fig)

Figure ?? plots the *geometric* average rate of return against the standard deviation of return for these asset classes. In general, riskier assets offered higher average rates of return — but only on average. Some, like commodities, ended up double trouble — they were risky *and* they left their holders worse off. Of course, when they originally bought commodities in 2005, investors did not know that this is what the future would hold. Such is the nature of risk. If investors had known what risk and reward they would be getting, they would probably not have purchased the commodities portfolio. (Then again, the nature of risk is that you don't know what will actually happen.) To get investors to bite, sellers' prices would have had to be lower, pushing up rates of return. And, of course, investors did not know. Similarly, they did not know in 2005 that stocks in general would have performed so well. (Sometimes, it's better to be lucky than smart.)

One glaring omission from Figure ?? are the rates of return on venture-capital (VC) and private-equity (PE) funds. These are high-risk, high-fee investment vehicles with long capital lock-up periods, nowadays sold by specialized advisors primarily to endowments. The prices by which they mark-to-market their investments are also

often more guesswork than precise. The investments cannot easily be transferred among investors. Even their past holding returns are difficult to come by. (In many cases, the endowment managers have to sign non-disclosure agreements to be “allowed” to participate. Perhaps not surprisingly, endowments have always performed poorly relative to the S&P 500. Hope never dies...) [Ang et al. \(JF 2018\)](#) have estimated the performance for the first four years in our sample. Relative to the S&P 500 index, they report

	Venture Capital				Private Equity			
	2005	2006	2007	2008	2005	2006	2007	2008
Median	-15%	-20%	-2%	-11%	+10%	+5%	-1%	-3%
Mean	+6%	-8%	+6%	+3%	+29%	+2%	+1%	+3%

These returns were after fees. In any case, these types of funds did not set the world on fire. Most finance professors are highly skeptical about investing in these funds. They are, however, very good business for their sponsors.

7.4 Individual Stocks

(Omitted fig)

Instead of buying entire asset classes, you could also have bought just one specific stock. Of course, I cannot show all 4,400 stocks by themselves, so I will just pick three. Figure ?? shows the rates of return on Intel Corp (INTC), Ford (F), and General Motors (GM). For comparison, the bottom is again the VFIAX (S&P 500).

You can see that individual stocks’ histograms are often really wide: Investing in a single stock would have been a rather risky venture, even though these were large companies and household names. Indeed, it is not even possible to plot the final year for GM in the rightmost compound return graph, because GM stock investors lost *all* invested money in the 2009 bankruptcy. (On the logarithmic scale, the final return was minus infinity.) Moreover, despite losing everything, the mean rate of return for a GM investor would have been only a modestly negative -4.5%. (You already know why: This was the difference between geometric and arithmetic averages explained on Page ??.)

(Omitted solvenow)

7.5 Comovement, Market Beta, and Correlation

Many important questions in finance are about the comovements among asset returns. Are the years in which the stock market did well also the years in which all its components and/or other assets did well, and vice-versa?

(Omitted fig)

Figure ?? plots a few return series. The X-axis is always the rate of return on the VFIAX (S&P 500) fund. The Y-axis is the rate of return on a second asset. Instead of points, the graphs name the relevant years (omitting the

'20' at the front). The diagonal line — with a slope of 1 — is dotted in black. The blue line is the best fit between the points. It is also called the “linear regression line.” Its slope — beta — plays a central role in finance.

The left column plots two individual stocks. It shows that when the **VFIAX** increased, Ford (**F**) increased even more *on average*. The slope of the blue line (called the market-beta) was greater than 1-to-1. This is primarily because Ford's stock more than tripled in 2009 when the stock market increased by 27%. (The large 2009 return on **F** also necessitated changing the scale on this panel.) In contrast, although Intel (**INTC**) also increased when **VFIAX** increased, Intel's stock did not increase *as much*. Its market-beta was less than 1.0. The most influential observation here was 2008, when both the stock market (–37%) and Intel's stock (–44%) crashed. But other observations were less clear. For example, in 2021, the market gained 29% while Intel gained only 6%. And both plots show noisy relationships. For example, Intel lost 15% in 2020, while the market gained 18%.

The market beta can also be more or less than 1.0 for industry funds. The middle column shows that the utility stock fund **VUIAX** had a slope of less than 1-to-1, while the information technology stock fund **VITAX** had a slope of more than 1-to-1. OK, I admit that I already hid the numbers in plain sight: If you go back to Table ??, it shows the market-betas of these funds, too. The betas of all stock-based funds tended to be around 1.0. Stocks with negative market-betas are rare. Even foreign stock markets tended to move together with the U.S. stock market at close to 1-to-1 market-betas.

The right column (and Table ??) show that the approximately 1-to-1 associations of **VFIAX** with other stocks did not extend to non-stock assets. For example, there was no clear association of rates of return between the stock market and Gold. Gold did not crash in 2008 (like **VFIAX**) but in 2013 (when **VFIAX** returned 32%). The long-term Treasury bond portfolio even had a modestly negative beta. This is partly because it rallied during the market crash of 2008 (as investors fled to the safety of government bonds) and dropped in 2013 (as investors gained more risk appetite). Was this foreseeable in 2005? Probably not. In the Covid years of 2020 and 2021, interest rates were so low that investors by-and-large preferred staying with stocks over government bonds.

Comovements are important if you are risk-averse. An asset that increased whenever the rest of your portfolio decreased was practically like “insurance.” It paid off when you needed it most. Put differently, this asset could have been valuable to you even if it offered only a low rate of return on average. In contrast, you would have been harmed by an asset that did worse whenever the rest of your portfolio also did badly. It would have increased your overall risk. You might not have liked this asset even if it had offered a higher rate of return on average.

Many investors hold primarily the overall stock-market portfolio, which is why they are particularly interested in covariation of other stocks with the market. The market-beta is a good measure of risk for adding a little bit of a new asset. There are three important cases:

- The slope was steeper than 1-to-1 (the diagonal). This implies that when the stock market did better (the X-axis), the asset did *a lot better* (the Y-axis) *on average*. Unfortunately, this would then also be the case when the stock market dropped. It would have dropped even more![‡]

For example, consider an investor who held primarily the stock market, and the stock market dropped by 10%. If an asset had a steep market-beta of 3.0, it would have lost 30%. On average, this asset made a bad situation worse. It added insult to injury.

[‡]In other time periods, we would have seen negative stock-market returns in many more poor years, allowing us to see the behavior of other assets in times when the overall stock market did poorly. Unfortunately, it is what it is. We will have to make do with these observations.

- The slope was between 0 and 1. For example, with a beta of 0.5, the 10% stock-market drop would have been associated with only a 5% drop in this asset. On average, this asset did not make the situation worse, but it also offered only modest consolation.
- The slope was below 0. For example, with a beta of -0.5 , this asset would have gained 5% in the 10% stock-market drop. On average, the asset was like insurance.

Instead of beta, you could measure the comovement with another statistic, the correlation. Correlation and beta are related. They always have the same sign. Yet the correlation has a feature that beta does not. A correlation of +100% indicates that two variables always perfectly move together; a correlation of 0% indicates that two variables move about independently; and a correlation of -100% indicates that two variables always perfectly move in opposite directions. (A correlation can never exceed +100% or -100% .) The correlation's limited range from -1 to $+1$ is both an advantage and a disadvantage. On the positive side, the correlation is a number that is often easier to judge than beta. On the negative side, the correlation has no concept of scale. It can be 100% even if one variable y moves only very, very mildly with x . For example, if a new fund consisted of 1% **VFIAX** and 99% **VMRXX** (the risk-free rate), this fund would have a perfect positive 100% correlation with **VFIAX**. But this fund really would have added very little risk. The market-beta correctly indicates this situation. It would have a market-beta of 0.01 — a low risk.

Of course, beta and correlation are only measures of *average* comovement: Even for assets with positive betas, there were individual years in which the asset and the stock market did not move together. For example, in 2020, the stock market did well, but **INTC** had a very bad year, losing 15% of its value. (Its competitor **AMD** produced better computer processors and doubled in value.)

(Omitted solvenow)

7.6 Will (150 Years of) History Repeat Itself?

Unlike a roulette wheel, where we know the statistical distribution of outcomes from the physical device, we do not know the distribution for most assets in the economy. Instead, we must assess future outcome distributions with the information we have.

For a financier, the best information about the future often comes from history. Of course, you are not interested in history for its own sake. It is useful only because it is your best available indicator of the future. Unfortunately, “best” is not particularly good. Fortunately, it is often good enough to be better than other investors. (In the land of the blind, the one-eyed man is king!) It is critically important that you understand the pitfalls of trusting history too much.

Inferences About Means

From 1880 to 2021, the overall arithmetic average rate of return on the stock market was 10.3% per year, the geometric average was 7.5% per year. The risk-free rate had an arithmetic and geometric average of 2.9% per year.

Let’s look at the much-longer-term historical performance of the S&P500 net of the short-term Treasury rate (both geometric) by decade (e.g., 1990s means Jan 1, 1990 to Dec 31, 1999):

1880s	1890s	1900s	1910s	1920s	1930s	1940s
2%	2%	6%	1%	11%	1%	8%
1950s	1960s	1970s	1980s	1990s	2000s	2010s
16%	4%	−1%	8%	13%	−4%	13%

If, in January 2000, you had believed that stocks would outperform Treasury bills, based on stocks’ 13% rate of return in their most recent 10 years, you would have been disappointed. The stock market underperformed short-term Treasuries by −4% per year instead. And fortunes turned again in the following decade. After the Great Recession of 2008-9, many finance professionals were gloomy and subsequently surprised by how well the market did in the 2010s.

And to me, the early 2020s were even more surprising. Despite Covid and the greatest recession since World-War II, the stock market started the 2020s with two bangs — returns of 18% in 2020 and 29% in 2021. The U.S. Treasury and the Federal Reserve had pushed interest rates down and propped up large firms (by extending cheap loans and other subsidies), so stocks outperformed fixed income. (Of course, at some point in the future, “we” — all of us — will have to pay for this. Rescues do not come for free.) Are similarly great returns likely to continue beyond 2021? Who knows?! Most academics do not expect the stock market to return 13% above Treasuries again — but they could be wrong.

Talking heads love to debate whether the fact that an asset had great returns in the past should now be interpreted to mean that the asset is also likely to be great in the future, or whether the asset is now overpriced (perhaps due to a bubble). Market prices are determined by a balance between optimistic and pessimistic views. Don’t look to me for guidance — I don’t know, either, and I am not a talking head. (I am only a writing head.)

Over longer horizons, could the future rate of return on the stock market be below that of Treasuries? Yes, but it seems unlikely. On the opposite side, could it be greater than 10% per annum? Yes, but it seems unlikely, too. We will return to this question in later chapters.

I can also tell you about a historical lesson. Investors in the Japanese stock market in 1986 saw the Nikkei-225 stock-market index rise from 10,000 to 40,000 by 1990 — a +40% rate of return per year. If they had believed that history was a good guide, they would have expected $40,000 \cdot 1.40^{13} \approx 3.2$ million by the end of 2002. Instead, the Nikkei fell below 8,000 in April 2003. It has only recently recovered to about 30,000 (as of late 2021). History would have been a terrible guide.

Similar puzzles apply to other risky assets. For example, real-estate returns (VGS LX) looked great in our sample from 2005-2021. However, many home owners lost their entire home equity (i.e., all their levered investments in their house) when house prices fell by more than 20% in 2008 and 2009. Are house prices too high now? I don’t

know. A good guess is that they are priced fairly, with expected returns in line with historical average rates of return and the risk they impose on their investors.

And then there is the possibility of a tail event — so named because it concerns a very low probability event far on the left tail (near –100%) of the outcome distribution. A popular finance writer, Taleb Nissim, dubbed a version of this a “black swan” event. I like to compare it to a large asteroid hit on Earth. How could you use the last 10,000 years of history to predict the (ultimately inevitable) incidence of another mass extinction caused by an asteroid impact? How can you use the last 100 years to predict a cataclysmic meltdown in the stock market?

I personally believe that crypto-currencies like Bitcoin or NFTs are collective Ponzi schemes. I see no value at the end of their chains. However, I cannot predict how these “assets” will do over the next few years. Will one bitcoin be worth nothing or \$1 million? The best predictions that I would venture are for large portfolios of today’s crypto-currencies over the next 20-30 years. (My prediction is that investors will likely lose a lot of money.) We will come back to these issues in Chapter

What about individual stocks? It is very difficult to assess expected returns even for large asset-class portfolios and long time-horizons. Attempting this for individual publicly-traded assets and over very short periods of time should be considered hopeless — your friends and even more animated talking heads on TV notwithstanding. (Caution: They also often tout their past investment successes based on selective memory.) Where active stock picking has worked over shorter horizons, researchers generally dismiss it as being more likely due to luck than skills. They recommend assessing only long-term trends, say over 1-5 years, and for large portfolios or asset classes. And even they do not know what the best length is. Having learned this the hard way, the professional investment industry is now generally moving from active management (where they try to predict individual stock performance) to passive management where they simply buy and hold larger portfolios or asset classes. This is also called factor investing.

A different way to pose the question is to ask when we are looking at a representative history and when we are looking at (just) an unusual time period. If this is not clear, make it even more extreme: In March 2020, Tesla (TSLA) traded for \$105/share. This was the onset of the Covid lockdown, when people would predictably soon drive less and all car manufacturers were pivoting towards electric cars in order to compete better against Tesla. Yet, by November 2021, TSLA traded for \$1,145/share — a 1.5 year rate of return of 900%! You should not imagine that its 900% historical return would be indicative of its expected future rate of return, either short-term or long-term. Just as Tesla backward-looking was just one random realization, Tesla forward-looking will be another. As difficult as it is to predict long-term average rates of return, predicting an individual year’s rate of return is even harder. Predicting the market’s stock return in any given year — or, worse, the stock return of just one stock — is really just a crapshoot. Similarly, do not expect General Motors to lose all stockholders’ money in the future because it did so in 2009 when it declared bankruptcy and all existing shares worthless.

A lottery analogy may help drive the points home. If you have played the lottery, your historical average rate of return is unlikely to be predictive of your future expected rate of return — especially if you have just won it big. Yes, you could trust history if you had millions of historical realizations, but you inevitably do not have so many. Consequently, your average historical payoff is only a mediocre predictor of your next week’s payoff. And you should definitely not trust your most recent realization(s) to be indicative of the future. Just because “5, 10, 12, 33,

34, 38” won last week does not mean that it will likely win again. It is only the long-run loss playing the lottery for many, many millennia that would be indicative of the long-run loss in future millennia (that is, unless the state sponsor changes the odds).

Inferences About Risks and Betas

Expected rates of return are very difficult to judge. They are fickle and unstable and notoriously difficult to predict. Interestingly, this is not the case for risk. Standard deviations and market-betas are more stable and easier to predict, especially for large portfolios.

From 1880 to 2021, the standard deviation on the stock market was about 16% per annum. With one exception — the Great Depression of the 1930s — the standard deviation on the stock market hovered in a narrow band between about 10% and 17% per year:

1880s	1890s	1900s	1910s	1920s	1930s	1940s
9.6%	11.8%	12.7%	10.5%	15.5%	37.7%	16.0%
1950s	1960s	1970s	1980s	1990s	2000s	2010s
11.9%	12.1%	15.9%	16.4%	13.4%	16.1%	12.5%

The numbers are almost the same if we subtract out the risk-free rate. I can confidently predict that the standard deviation will be between 12% and 20% per year over the next 30 years — even though I cannot predict whether the stock market will return –5% per year or +10% per year *on average*.

Let me also give you a preview for practical estimation advice. Future standard deviations and market-betas are better predicted from higher-frequency data than from annual rates of return. A good choice is to use about 1-3 years of recent daily rates of returns. Furthermore, it is also useful to “shrink” the resulting estimates towards the overall average statistic for all assets. For example, if you estimated the market beta of a particular stock to be 2.0, you should average it with 1.0 (which is the average market-beta in the stock-market) to predict a future market-beta for this stock to be about 1.5. There are even better estimation methods, but better beta estimation is better covered in an advanced financial econometrics course.

Summary

Here are the two most important lesson about history to remember:

- Risk:** Historical standard deviation and market-betas are reasonably reliable predictors of future standard deviations and market-betas, especially for diversified large portfolios over reasonably long time spans (and of course assuming good estimation procedures).
- Reward:** Historical average rates of return, especially of individual assets and over shorter time spans, are not very reliable predictors of future expected rates of return.

With these caveats, henceforth, as in almost all of finance, we will just assume we know the statistical distributions from which future investment returns will be drawn. For exposition, this makes our life a lot easier. When you want to use our techniques in the real world, you will usually collect historical data and pretend that the future

However, always remember: historical data is an imperfect guide to the future. Trust your risk estimates more than your expected return estimates.

(Omitted solvenow)

7.7 Conclusion: The Big Picture

What could you learn from our little excursion in this chapter? Actually, almost everything there is to learn about investments, though they will require more detail that will come in the next chapters. So here are the most important points summarized. First, risk:

- Risk is often fairly stable and predictable.
- Cash is the safest investment — its distribution is tightly centered around its mean. Bonds are riskier. And stocks are even riskier. (Sometimes, stocks are said to be “noisy” — their standard deviation is so high that it is exceedingly difficult to predict how they will perform.)
- You can lose your shirt investing in individual stocks, even prominent ones.
- Large portfolios consisting of many stocks tend to have lower risk than their individual stock components. (This is due to “diversification,” a concept we will discuss in the next chapter.)

Second, comovement:

- Comovement is often fairly stable and predictable.
- Most stocks tend to move together. Bonds, cash, and alternative investments tend to have lower correlations with stocks.
- The fact that some investment rates of return tend to move together more than others is important. It is the foundation for beta, a measure of risk that we will explore in Chapter . Beta determines how effective diversification can be.

Third, reward:

- Expected returns are not easily predictable and many economists believe that they are not greatly stable, either. History can be a poor guide to their future.
- Stocks have offered higher average rates of return than bonds, which in turn have offered higher average rates of return than cash. However, keep in mind that this was only *on average*. In any given year, this might have been the opposite. And it may not necessarily be the case even over coming decades.
- There was a general relationship between risk and reward: Riskier investments tended to have higher average rates of return. (However, you will also learn soon that the appropriate risk measure depends on the context. Thus, please do not overread the simple relationship between the mean and the standard deviation here.)
- The arithmetic (average) rate of return is always larger than the geometric (compound) rate of return. A positive average rate of return usually, but not always, translates into a positive compound holding rate of return. Indeed, you can lose all your money even when the arithmetic average rate of return is positive.

A Appendix: 150 Years of Monthly Stock-Market Returns

(Omitted fig)

Figure ?? shows a histogram of about 150 years of monthly rates of return on the stock market (an S&P 500-like portfolio including dividends) with an overlaid density plot similar to that used in the main text of this chapter. This graph shows (1) the performance over a longer time span; and (2) the fact that histograms of stock-market returns are usually *not* bimodal. The average rate of return was 0.86% per month for the index (with dividends), 0.56% net of the risk-free rate, both with a standard deviation of about 4.7% per month. The geometric averages were 0.75% and 0.45% per month, respectively.

B Appendix: Causality vs. Correlation

Now that you understand the concept of correlation, allow me a brief diversion. The most important problems in finance — and economics, and statistics, and science, and theory, and practice — may well be whether correlation implies causation. If X causes Y, the two should be correlated. The problem is that if X does not cause Y, the two can still be correlated. People do not have very good intuition about the distinction. Thus, they often commit serious and harmful interpretation mistakes. (Some are deliberate, as in political demagoguery.)

For example, just because a stock had tended to go up when there was sunspot activity (or when the overall stock market went up) does not mean that you have found that sunspots are one cause of stocks going up — this is an example of spurious correlation. (In fact, over long sample periods, increased sunspot activity has indeed been associated with higher stock prices.) Just because good CEOs are paid high salaries on average does not mean that paying more money would make your own CEO any better — the causation may go the other way. And just because increases in government spending are associated with reductions in unemployment does not mean that the government can reduce unemployment by spending more — some other economic factor may have determined both spending and unemployment. Determining causality is important if you want to know how your strategies and policies are likely to change outcomes.

An early answer to measuring causality in economics came from two econometricians: if unexpected changes in X predict unexpected changes in Y, then X may cause Y. If you saw an unusual sunspot (X) and it usually precedes (in time) an unusual increase in the stock return (Y), then it is a hint that X causes Y. This concept is called Granger-Sims causality. By this metric, the data reject the hypothesis that sunspots have “caused” stock returns to go up. (Both sunspots and stock prices happened to go up; it was time effects that induced the spurious correlation.) Unfortunately, Granger-Sims causality isn’t perfect, either. By its metric, the weather forecast “causes” the weather. (Unusual changes in the forecast indeed predict subsequent unusual changes in weather!)

Just when we were ready to give up, economics stumbled upon an approach that is now called “quasi-experimental” (and won the 2021 Nobel Prize in economics). It is revolutionizing empirical economics right now (and soon economic consulting, too). Let me illustrate this with an example. Think about figuring out whether access to loans increases the success of startups. The problem is that ventures that are less likely to be successful are also less likely to attract lenders. Thus, it is not possible to conclude from the fact that funded ventures had higher success rates in the past (which they did!) that loan access played a critical role in this success. Funding may have been more like the weather forecast, itself responding to other factors (e.g., promising business plans) that ultimately determined project success. If loans to startups did not have a positive influence on survival, government programs that seek to make more loans available to more startups would probably not be a good idea.

But [Fracassi, Garmaise, Kogan, and Natividad](#) have an answer! It turns out that a particular lender employed an automated credit score algorithm with a cutoff that determined loan funding. Applicants with scores just above the cutoff (say, 4.14) received a loan. Applicants with scores below (say, 4.13) were denied. The probability of survival was 30% for the 4.14 group and

25% for the 4.13 group. Because these applicants were so close in score, their differences were probably just noise. Thus, it is likely that the entire 5% was due to the fact that the 4.14 firm got the loan and the 4.13 firm did not. (To make sure, they also compared their 5% survival difference to the survival differences for firms between 4.12 and 4.13 and firms between 4.14 and 4.15, neither of which was treated differently by the lender and neither of which showed any differences in survival.) This is convincing evidence that access to loans indeed helped improve the chances of survival for the startup firms *in their sample*.

C Appendix: Equity-Related Market Institutions

Although there is more capital deployed in non-equity than equity financial markets, stocks remain the most interesting financial instrument from a corporate perspective. Stocks finance most of the risks of corporate projects. Moreover, the subject area of investments has traditionally focused on equities (stocks), historical data for stocks is relatively easy to come by, and stocks are among the most accessible investments to retail investors. So it makes sense for us to describe more of their institutional details.

In 2021, there were about [4-5,000](#) publicly traded stocks in the U.S. with a market capitalization of about [\\$50 trillion](#) — more than twice the total U.S. gross domestic product of \$22 trillion. That's about \$65,000 per American (though wealth and stock holdings are quite unequal, and a lot of the shares are held by foreigners). New companies raised about [\\$500 billion](#) in IPOs from U.S. investors, and publicly traded companies raised another \$300 billion. Between \$10 and \$15 billion worth of equities changed hands on an average day. These impressive statistics merit a peek behind the mechanics of the stock market in more detail.

Brokers

Most individuals place their orders to buy or sell stocks with a retail broker. The landscape of retail brokerage services has been changing greatly over the decades. It now includes older institutions (such as [Fidelity](#), [JP Morgan](#), [Merrill Lynch](#), [Charles Schwab](#), and [Vanguard](#)) that predate the Internet and offer human investment advisors; companies that entered the business when the Internet became popular (such as [E-Trade](#), [Interactive Brokers](#), and [TD Ameritrade](#)); and the latest entry, [Robinhood](#), which caters to hyper-small retail investors.

Retail brokers have two functions. First, they handle the execution of trades. They usually do so by routing their investors' orders to a "wholesaler" or a centralized trading location (e.g., a particular stock exchange, such as the New York Stock Exchange [NYSE], the "National Association of Securities Dealers Automated Quotation System" [Nasdaq], or a large market-maker [like Citadel Securities]). Second, they keep track of investors' holdings, facilitate buying on margin (whereby investors can borrow money to buy more stocks), and help in short selling of securities, by which investors can speculate that a stock will go down and which will soon be explained below.

Many large institutional investors separate the two functions: They employ their own traders, while the broker takes care only of the bookkeeping of the investor's portfolio, margin provisions (lending money to the investor against the collateral of the stocks themselves), and shorting provisions. Such limited large-scale brokers are called prime broker[prime brokers].

(Omitted solvenow)

Orders and Wholesalers

Investors can place different types of orders. The two most common are the market order, which requests execution at the current price, and the limit order, which requests execution only if the price is above or below a limit that the investor can specify. (There are also other types, such as *stop-loss* orders [which instruct a broker to sell a security if it has lost a certain amount of money], *good-till-cancelled* orders, and *fill-or-kill* orders.)

What happens when your retail broker sends your market order to a wholesaler? During standard hours, the wholesaler practically guarantees instant execution. If the market-maker does not have another customer to match your order to, the market-maker becomes the counterparty itself (sort of like a used-car dealer who holds inventory). Holding inventory can, at times (e.g., in a market crash), become risky business. However, the intermediaries typically also have access to more information, such as how many investor limit orders are standing by to buy and sell shares if the price changes, which generally limits their risk.

Some less active venues (e.g., stock markets in smaller countries) have organized financial markets into noncontinuous auction systems, which match buy and sell orders only a couple of times each day. The disadvantage is that traders cannot execute orders immediately but have to wait until a whole range of buy and sell orders have accumulated. The advantage is that this arrangement eliminates the risk that an (expensive) intermediary would otherwise have to bear. Thus, auctions generally offer lower trading costs but slower execution. Even in the United States, large block holders (owners of large blocks of a company's shares or bonds) often participate in specialized auction markets (Archipelago, Instinet, or Posit) rather than in the standard continuous markets.

Although computerization has made many financial markets more perfect and liquid (with high transaction volumes and low transaction costs), this is not the case for all securities. There are still many securities that are traded over the counter, i.e., by traders calling up one another over the telephone or an instant messaging system, and negotiating a price. Importantly, other parties may not learn the price at which the securities traded. It is unwise for greenhorns to participate in such markets — the traders will quickly figure out the expertise of who they are talking to and take advantage of their better information.

(Omitted solvenow)

Execution Costs

Robinhood's aggressive pricing approach has pretty much forced all retail brokers to go to a model where Internet-placed orders no longer charge fees. Without fees, how do retail brokerage firms then make money? The system is a little surreptitious. Their primary source of income is that they receive "rebate" payments back from the wholesaler to whom they route the order. This arrangement is called payment for order flow. The wholesaler in turn recoups this payment to the broker by executing your trade at a price that is slightly less favorable to you. The difference between the buying and selling price is the famous bid-ask spread. For example, a stock may have a bid of \$49.99 and an ask of \$50.00. If you place a small buy order, you pay \$50.00; if you want to sell, you get \$49.99. I chose these numbers deliberately: the typical bid-ask spread nowadays is indeed 1-3 cents per share, guaranteed for small orders. The Securities and Exchange Commission (SEC), a federal agency established in 1933/1934 to protect investors, has enacted regulations that limit wholesaler execution to be no worse than what large stock exchanges are posting. Unfortunately, it is also almost impossible to determine which brokers offer the best execution — there are federal laws on the books that have the unintended side-effect of making it illegal to probe execution quality.

In sum, despite the bid-ask source of friction, trades are still fairly cheap, but just not as cheap as they appear. Round-trip transactions still cost investors money, and market-making is still a very profitable business at large scale.

The result is that with no fixed fee and this low a spread, even small investors can now “day-trade” small numbers of shares (and many do) without losing a great amount of capital in transaction costs. Nevertheless, day trading is rarely a smart strategy. Many day traders eventually lose a good deal of money — similar to many gamblers in casinos.

High-frequency traders (HFTs) are typically large speculators that trade even more than day traders. The biggest HFTs spend a lot of money to co-locate their trading servers at exchange computers in order to shave a nanosecond’s advantage over their peers. The exchanges profit through higher trading volume, i.e., in turn benefitting from the bid-ask spread, so they welcome HFTs.

HFTs place and withdraw millions of orders every day. Your first question should be “How can HFTs make money in a competitive market?” The answer is that their business is trying to front-run the rest of us (and nowadays one another, too). Front-running means that if any trader has posted a limit order and does not revise it quickly enough when new information arrives, then the fastest HFT who has learned that the price has changed can transact at (pick off) their limit offers. The reason is that markets are not really competitive in nanoseconds. (Light travels only 20cm in a nanosecond.) There is only one HFT at the very nanosecond it pounces on a standing limit order, even if another trader would have loved to pay a tiny bit more another nanosecond later. Thus, there has been an arms race among HFTs to be faster than others — and the speed of light has literally become the constraint!

There is some debate about whether HFTs add liquidity to the financial markets or siphon it off. Because their activity is anonymous and thus dark to independent academic research, we can only guess. Some academics have suggested bunching orders into auctions once every second, but it is hard to know what the optimum is here (1 millisecond? 1 second? 1 hour?). But this is probably no longer a big problem. HFT was very profitable in the beginning. However, over time, as in any other business, competition has eroded HFT profits. Competition cures a lot of ills. It is no longer clear whether HFTs earn profit by taking more risk or by arbitrage. Besides, even if HFTs are a problem, it is never clear whether government intervention would improve or worsen the situation. (If you are intrigued now, the popular author Michael Lewis has written a popular and critical [book](#) on the subject.)

(Omitted solvenow)

How Shorting Stocks Works

If you want to speculate that a stock will go down, you would want to short it. Perfect markets assume that it is costless to do so. In the real-world, the process is more nuanced and requires more explanation.

Shorting is arranged by stock brokers. Here is how it works:

- The broker finds an investor in the market who is willing to lend you the shares. In a perfect market, this transaction does not cost anything. In the real world, it does. Both the broker and lender usually earn a few basis points per year for doing you the favor of facilitating your short sale. Many large index funds earn most of their money by lending out shares rather than from fees from their own fund investors.
- After you have borrowed the shares, you sell them into the market to someone else who wants to buy the shares. In a perfect market, you would keep the proceeds and invest them any way you wish. You could earn interest on them. In the real world, your broker may force you to put these proceeds into low-yield safe bonds. If you are a small retail investor, your brokerage firm may even keep the interest proceeds altogether.
- When you want to “unwind” your short, you repurchase the shares and return them to your lender.

For example, if you borrowed the shares when they were trading for \$50 (and sold them into the market), and the shares now sell for \$30, you can repurchase them for \$20 less than what you first sold them for into the market. This \$20 is your profit. In an ideal, perfect world, you can think of your role effectively as the same as that of the company — you can issue shares and use the \$50 proceeds to fund your investments (e.g., to earn interest). In the real world, you have to take the various transaction costs into account.

Shorting has become so popular that there are now other ways to do it, too. Shorting bundles of stocks, such as the stocks in the [S&P 500](#) index or in some other market indices (and even some particularly popular stocks) has become particularly easy. For example, you can short the relevant index exchange-traded fund (an ETF, explained below, which works the same way as shorting any other stock); or you can sell traded Futures (which are like side bets between interested buyers and sellers) on common stock-market indexes. Shorting is also common and easy for bonds.

(Omitted solvenow)

How Securities Appear and Disappear

► Inflows

Most publicly traded equities appear on public exchanges, most often Nasdaq, through initial public offering[initial public offerings] (IPO[IPOs]). This is an event in which a privately traded company first sells shares to ordinary retail and institutional investors. IPOs are usually executed by underwriter[underwriters] (investment bankers such as Goldman Sachs or Bank of America's Merrill Lynch), which are familiar with the complex legal and regulatory process and have easy access to an investor client base to buy the newly issued shares. Shares in IPOs are typically sold at a fixed price — and for about 10% below the price at which they are likely to trade on the first day of after-market open trading. (Many IPO shares are allocated to the brokerage firm's favorite customers, for whom they can be an important source of profit.)

(Omitted anecdote)

Usually, about a third of the company is sold in the IPO, and the typical IPO offers shares worth between \$20 million and \$100 million, although some are much larger (e.g., privatizations of previously national companies, or Facebook's \$16 billion IPO in 2012). About two-thirds of all such IPO companies never amount to much or even die within a couple of years, but the remaining third soon thereafter offer more shares in seasoned equity offering[seasoned equity offerings] (SEO[SEOs]). These days, however, much expansion in the number of shares in publicly traded companies — especially for large companies — comes not from seasoned equity offerings but from employee stock option plans, which eventually turn into unrestricted publicly traded shares.

The SEC is also in charge of regulating certain behaviors of publicly traded companies. This includes how they conduct their IPOs. It also describes how they have to behave thereafter. For example, publicly traded companies must regularly report their financials and some other information. Moreover, Congress has banned insider trading on unreleased *specific* information, although more general informed trading by insiders is legal (under what is called a 10b-5 plan and which seems to have been fairly common and profitable). Moreover, there are other loopholes that allow smart CEOs and politicians to trade legally on inside information. (These do not apply to funds and external investors.) The SEC can only pursue civil fines. If there is fraud involved, then it is up to the states to pursue criminal sanctions, which they often do simultaneously. (In addition to SEC rules, publicly traded firms have to follow a hodgepodge of other federal and state laws.)

Because IPOs face unusually complex legal regulations and liability, the alternative of reverse merger[reverse mergers] has recently become prominent. A larger privately-owned company simply merges with a small company (possibly just a shell

company with no other assets) that is already publicly traded. The owners of the large private company receive newly issued shares in the combined entity. And, of course, the newly issued shares in effect move private-sector assets into the public markets, where the firms' assets then appear in the form of additional market capitalization. A recent popular variant of this way of going public was to found a Special Purpose Acquisition Company (SPAC) first, which is the mostly empty shell company that goes public. With only cash, SEC regulations regarding the company's activities are easy to satisfy for the SPAC. Once public, the sponsors of the SPAC then look for private firms to purchase.

► Outflows

Capital flows out of the financial markets in a number of ways. The most important venues are capital distributions such as dividends and share repurchases. Many companies pay some of their earnings in dividends to investors. Dividends, of course, do not fall like manna from heaven. For example, a firm worth \$100,000 may pay \$1,000, and would therefore be worth \$99,000 after the dividend distribution. If you own a share of \$100, you would own (roughly) \$99 in stock and \$1 in dividends after the payment — still \$100 in total, no better or worse. (If you have to pay some taxes on dividend receipts, you might come out for the worse.) Alternatively, firms may reduce their outstanding shares by paying out earnings in share repurchases. For example, the firm may dedicate the \$1,000 to share repurchases, and you could ask the firm to use \$100 thereof to repurchase your share. But even if you hold onto your share, you have not lost anything. Previously, you owned $\$100/\$100,000 = 0.1\%$ of a \$100,000 company, for a net of \$100. Now, you will own $\$100/\$99,000 \approx 0.10101\%$ of a \$99,000 company — multiply this to find that your share is still worth \$100. In either case, the value of outstanding public equity in the firm has shrunk from \$100,000 to \$99,000. We will discuss dividends and share repurchases in Chapter .

Firms can also exit the public financial markets entirely by delisting. Delistings usually occur either when a firm is bought by another firm or when it runs into financial difficulties so bad that it fails to continue meeting minimum listing requirements. Often, such financial difficulties lead to Chapter 11 bankruptcy or Chapter 7 liquidation (so named for their chapters in the U.S. Bankruptcy Code). Some firms even voluntarily liquidate, determining that they can pay their shareholders more if they sell their assets and return the money to them. This is rare because managers usually like to keep their jobs — even if continuation of the company is not in the interest of shareholders. Unfortunately, bankruptcy usually means that equity investors lose all their investment. Fortunately, equity investors also enjoy limited liability, which means that they can at most lose their original investment. They do not have to pay further for the sins of the firm with their other personal assets (unlike partners, such as investor-like “names” in Lloyd's of London, who are on the hook with everything they own personally if they participate in an insurance underwriting venture).

(Omitted solvenow)

Investment Companies and Vehicles

In addition to its regulation of how companies go public and what insiders may trade, the SEC also regulates investment advisors and funds. In practice, this has allowed three different types of regulated investment companies (fund[funds]) to operate in the public markets: (1) open-end funds, (2) closed-end funds, and (3) unit investment trusts (UITs).

A fund that is open-end must be ready to repurchase its shares (or hand the requesting investor the underlying [other] company shares represented by the fund shares) when its investors demand it. Most open-end funds not only repurchase but also sell new shares in themselves at the same price. There are two common types of open-end funds that are close alternatives from the perspective of investors: mutual funds and exchange-traded funds. A mutual fund must post its net asset value (NAV), at the end of each trading day. (Thus, mutual funds do not trade on exchanges.) This NAV is effectively the price at which

investors can purchase or redeem mutual fund shares. An exchange-traded fund[exchange-traded funds] (ETF[ETFs]) has different rules about who is allowed to ask for new shares to be created or existed shares to be removed. Its shares trades just like ordinary company shares on stock exchanges. There are now about 8,000 open-end mutual funds with \$25 trillion in assets and 3,000 exchange-traded funds with \$6 trillion in assets in the United States alone. Note how there are now many more stock funds than individual stocks.

The redemption right gives the law of one price 2.5]subject:lawofonepriceLaw of One Price a lot of bite — fund shares are almost always worth *nearly* exactly what their underlying holdings are worth. If an open-end fund's share price were to fall much below the value of its holdings, an arbitrageur could buy up the fund shares, redeem them, and thereby earn free money. (One discrepancy is due to some odd tax complications: the fund's capital gains and losses are passed through to the fund investors at the end of every year, but they may not be what every investor experienced.)

In a closed-end fund, investors cannot demand to redeem their fund shares for the underlying value. The advantage of a closed-end fund is that it can itself invest in assets that are less liquid (such as a building). After all, its investors cannot force a closed-end fund to sell its holdings suddenly based on a whim. (Many closed-end funds are exchange-traded, so that if a closed-end fund investor needs cash, she can resell her shares.) The disadvantage of the closed-end scheme is that the law of one price has much less bite. On average, closed-end funds trade persistently below the value of their underlying holdings, roughly in line with the (often high) fees that the managers of many of these closed-end funds are charging.

Both mutual funds and closed-end fund managers are allowed to trade fund holdings quite actively — and many do so. Although some funds specialize in imitating common stock-market indexes, many more try to guess the markets or try to be more “boutique” (specialized and unusual). Most funds are classified into a category based on their general trading motivation (such as “market timing,” or “growth” or “value,” or “income” or “capital appreciation”).

A unit investment trust (UIT) is sort of closed-end in its creation (usually through one big primary offering) and sort of open-end in its redemption policies (usually accepting investor redemption requests on demand). Moreover, regulatory rules forbid UITs to trade actively (although this is about to change), and UITs must have a fixed termination date (even if it is 50 years in the future). UITs can be listed on a stock exchange, which makes it easy for retail investors to buy and sell them. Though there are a few thousand UITs in the United States, their total asset value is only about \$0.1 trillion.

Some other investment vehicles are regulated by the SEC under different rules. The most prominent may be the American Depositary Receipt (ADR). An ADR is a passive investment vehicle that usually owns the stock of only one foreign security, held in escrow at a U.S. bank (usually the Bank of New York). ADRs make it easier for U.S. retail investors to trade in foreign securities without incurring large transaction costs. ADRs are redeemable, which gives the law of one price great bite.

There are also funds that are structured so that they do not need to register with the SEC. This means that they cannot openly advertise for new investors and are limited to fewer than 100 investors. This includes most hedge fund[hedge funds], venture capital funds, and other private equity funds. Many offshore fund[offshore funds] are set up to allow foreign investors to hold U.S. stocks not only without SEC regulation, but also without ever having to tread into the domain of the U.S. tax authorities.

(Omitted solvenow)

Summary

This chapter covered the following major points:

- It provided a summary of asset-class performance from 2005 to 2021:

- Stocks, on average, had higher average rates of return than long-term bonds, which in turn had higher average rates of return than cash investments.
- Individual stocks were riskiest. (Large-firm-type) stock-market portfolios had lower risk than individual stock holdings. Bonds had modestly lower risk. Cash was least risky.
- Stocks have outperformed bonds by about 6% per year and cash by about 10% per year.
- Not all riskier asset-classes delivered higher rates of return.

Over much longer horizons, stocks outperformed long-term bonds by about 4% per annum and cash by about 6% per annum.

- Stocks (and many other investments) have tended to correlate positively with one another: When the stock market overall has had a good year, most individual stocks have also tended to have a good year (and vice-versa). Long-term bonds have tended to correlate negatively with the stock market over the last few decades (but not before).
- Most finance assumes that investors know such statistics as means, standard deviations, and betas — as if you understood the mechanical device generating the randomness. This is a leap of faith. In real life, historical data can help you in predicting the future, but it is not perfect. Historical risks and correlations are often good predictors of their future equivalents, but historical means are not.
- There is an important distinction between causality and correlation. Economists have recently developed some good causality tests.
- Section ?? explained many institutional arrangements governing publicly traded equity securities. This includes the roles of retail and prime brokers, exchanges, and funds. It also described how stocks can be shorted, and how funds flow in and out of the financial markets.