

PEABODY RIVER ASSET MANAGEMENT, LLC

P.O. Box 382128
Harvard Square
Cambridge, Massachusetts 02238-2128
www.peabodyriver.com

NEWSLETTER

Issue 11

April 2011

NEWS OF PEABODY RIVER

On the last day of the first quarter of 2011, Peabody River Asset Management doubled in size. Robert A. Budding joined the firm as a portfolio manager. Last year, a friend skeptically asked me, when I used the word “we” in reference to Peabody River, whether there was someone else working here. I explained that I used the corporate “we” to refer to the firm, and “I” to refer to myself. And now, Peabody River really is a “we” in every sense.

Bob was a classmate of mine at the Booth School of Business at the University of Chicago. Although Bob’s professional career to date has been in corporate financial management--mergers and acquisitions, risk management, and strategic planning--his real interests have been in macroeconomic analysis and investment management. For some time, he has offered free investment guidance to friends, and he and I have for years discussed investing over the occasional coffee. His analytical framework is the same as mine, so our combination is a perfect fit.

From the instant he joined Peabody River, Bob officially took on a couple of clients, friends who have known him a long time and have sought his financial advice. And they have now referred others to him.

I should make clear that Bob will have his clients, and I will have mine. Although we talk through problems together, there is only one portfolio manager, Bob or I, serving each client. But the reason I invited Bob to join me was that I wanted to give Peabody River’s clients, and its prospective clients, the assurance that the firm has depth, and that if, for any reason, I am unavailable, there is someone else whom they can call. Of course, our clients’ assets have always been safe and accessible regardless of whatever may befall us, because Peabody River follows the best practice of our profession by leaving them in the custody of a third party.

Bob is a native of California. He was graduated from the University of Connecticut *summa cum laude* in chemistry and was pursuing a doctorate in physical chemistry (my late father’s field) at Dartmouth College when he decided to change course, and to make his career in business, which is how we

came to meet in Chicago. As most of Peabody River's friends know, my academic background is in history, and Bob has a connection to that discipline, too, through marriage; his wife, Audrey, has a doctorate from Harvard in Balkan history, and she now teaches history at a private school in Boston.

We're still working on our integration, so that our telephones and computers will be linked, and Bob's portrait and biography will appear on the Peabody River Web site. This will happen over the next couple of months.

I am delighted to welcome Bob to the company without waiting for these technicalities to be completed.

§

BRIEF REVIEW OF THE FIRST QUARTER OF 2011

The stock market confounded the bears in the first quarter of the year. The financial press was full of caution at the start of the year, and it still is, because, by conventional valuation measures, the market didn't seem to offer much possibility for more return. At any rate, the total return on the S&P 500 was 5.92% for the three months. This simply reinforces our commitment not to place our trust in short-term market forecasts. As some of the press noted, this was the best first quarter since 1998. I'm not sure why that matters, but it gives a little perspective to the number. In the performance review in my previous newsletter, I mentioned that it was entirely possible for the stocks of large companies, at least for a while, to outperform the stocks of smaller companies. That wasn't a forecast—I always emphasize that I don't make short-term forecasts—but it might as well have been: The stocks of smaller companies had a return of only 0.5% (as measured by the S&P 600) during the quarter.

One measure of the world's stock markets, excluding the U.S., was up 3.41% (in dollars), and the emerging markets were up only 2.05%.

The total return for the U.S. bond market (as represented by the Barclays Capital U.S. Aggregate Bond Index) was 0.42% for the same period.

It's been a boom time for commodities, which were up 4.45% (as measured by the Dow Jones-UBS Commodity Index, which is a broad measure that we use for our investing).

ESSAY: WHAT RETURN CAN WE EXPECT FROM STOCKS?

The race is not always to the swift nor the battle to the strong—but that’s the way to bet.

Damon Runyon

Part 1: The Question

A friend of mine who was studying the philosophy of science at Cambridge University told me that an eminent don in his department assumed that every question was necessarily a profound one. If you asked him the simplest question, he would ponder awhile, and then slowly begin, “I think I see what you’re getting at...”

Some of my readers who thought that they already understood the basics of investing likely feel as if they’re in the presence of that don when they encounter my essays. But I’m not imposing complications on investment questions where none would otherwise exist. In investing, clarity of thought requires the careful dissection of concepts that are too often taken for obvious. And so, this essay begins with the exegesis of a seemingly simple, innocent question:

What return can we expect from stocks over the long term?

This sentence contains four problematic terms: “return,” “expect,” “stocks,” and “long term.”

Let’s consider each in turn. Only once we understand these terms can we begin to formulate an answer. The question itself is critical; without an answer, we’d be hard pressed to justify investing in stocks at all. True, the lack of a view on this doesn’t stop many investors from heedlessly risking their money in the market. But why would you invest if you didn’t have a view on whether, say, stocks were likely to have a higher return than cash?

1. **Return.** We’ve already covered the meanings of “return” in an earlier essay.¹ In the context of this question, it means the **total return**, which is the sum of the return from income (the dividend that a stock pays out) plus the return from the change in price, that is, the capital gain or loss. Also in this context, it means the annual **rate of (total) return**. It doesn’t mean the cumulative change, as when I bought a painting for \$100,000 in 1997 and sold it for \$150,000 in 2007. That’s a *cumulative* return of 50% over ten years, but the corresponding *rate* of return is just 4.1% per year.

If you multiply that 4.1% by 10 years, you won’t get 50%. That’s because the rate is a *growth* rate, and growth begets new dollars which beget even more growth. This is the essence of compounding, which we considered in another essay.² Because of compounding, a small positive rate can produce a large cumulative return over time; in other words, the cumulative

¹ Peabody River *Newsletter*, issue 2, July 2008, “[How to Think about Returns.](#)”

² Peabody River *Newsletter*, issue 7, April 2010, “[How Much is that Investment Worth in Real Money?](#)”

return (if positive) is always greater, sometimes much greater, than the rate of return times the number of years.

Furthermore, the *average* of the one-year returns is not the same as the annual *rate* of return, though they both summarize the one-year returns. This is a subtle but significant distinction that is missed by almost everyone unfamiliar with the mathematics of investing. Unless the one-year returns are all the same, their average will always and necessarily be greater than the rate. To begin to see why, consider the fanciful example of the stock market dropping 50% in the first year, then going up 100% in the second year. The net result is no change in value (because your investment halved, then doubled), for an annual *rate* of return of 0%. But the *average* return was $\frac{-50\% + 100\%}{2} = 25\%$, not 0%.

We'll spend most of this essay considering the annual *rate* of return. For brevity, I'll sometimes write just "return" when I mean "rate of return," but the context should make the meaning clear. Toward the end, we'll consider the average annual return, because this is what you might think of as a typical return.

- 2. Expect.** What we expect to happen seldom does happen, at least, not exactly as we expected it. But—I should hope—we don't live our lives without regard to the future; rather, we act, and sometimes even plan, according to some sort of expectations. "Expectation" or **expected value** has a precise meaning in the discipline of statistical analysis. It does not mean what we hope or fear will happen (at least, not necessarily). It also, more subtly, does not mean the one result from among all the possibilities that is most likely to happen. Rather, it denotes the average of the *possibilities* we see before us, but that average is adjusted to reflect their *probabilities*. The more likely results contribute more to the average.

The expected return is not at all the same as the average return I described a moment ago. That was the average of the year-by-year returns from the present into the future. The expected return, rather, is the average of all the possible rates of return that may have been realized between now and a specified time in the future, weighted according to their likelihood.

That is to say that, although when we arrive in the future and look back, we will see only one sequence of year-by-year returns from our stock market investment, we don't yet know what that sequence will be, and there is an infinitude of future possibilities. From our current vantage point, there are possible sequences that could result in very high rates of return; there are possible sequences that could result in very low rates of return. Both extremes are unlikely. Much more likely are the possible sequences of year-by-year returns that result in middling rates of return. The expected return is the average of all these possible rates, but biased toward the rates that correspond to the more likely future market developments.

Don't marvel that to come up with our expected value, we'll have to calculate the infinite number of possible sequences of returns and their probabilities in order to calculate the average of the resulting rates. We won't, but the concept of expected value as an average will prove useful when we have to situate the expected return in the context of risk, as any reader of my essays will know that we inevitably must.

- 3. Stocks.** When I say “stocks,” I mean the *entire* stock market (for the purpose of this essay, the U.S. stock market). You can expect a selection of stocks to have a different return from that of the market as a whole. If *you* have only five stocks in your portfolio, the answer to our question won’t be the same as the answer to, “What return can we expect from *your* stocks over the long term?” “Stocks” as the entire market is a sort of abstraction, an ideal rather than a practical investment. But for the last few decades, you could, if you so wished, buy something very close to the entire market by buying into a mutual fund designed to replicate the S&P 500 or some other index intended to represent the entire U.S. stock market, and more recently, there have been exchange-traded funds (ETFs) that allow you to do the same thing. There are costs associated with these investments that reduce the return in practice below the return in theory, but these costs are small, and so the differences between the practical and the theoretical returns are very small. In answering the question, we’ll assume away all costs.

“Stocks,” the ones that constitute the entire market, are not a static group over time. Companies have lives; they are founded, grow, become listed on stock exchanges, shrink, are taken private, or merge out of existence. You might hold a portfolio of a few stocks that, if you’re lucky, will continue to exist over long spans of time. But the market as whole, in its composition, is in constant flux, as are the indices that track the market.

Because we are considering total return, we have to assume that all the dividends you receive are immediately reinvested in the stock market. This is another way of saying, once again, that our question concerns the **total return**. Our question, as I’ve framed it, does not address the practical question of what would happen to your investment in the stock market if you were to live off the dividends, and you wanted to know the return on the stocks that were left behind. If you were to devour the dividends rather than to reinvest them so that they could themselves earn returns in the future, the resulting market return would be much less than the answer we’re seeking. It’s not a bad practical question if you want to know how to live off your investments, but it’s worth asking only after we know how much total return the market is capable of providing.

- 4. Long term.** Some uneducated pundits who hold forth on one of the pressing issues of our time, global climate change, are unable to distinguish between forecasting the climate and forecasting the weather, and therefore treat both as equally uncertain, which they’re not. As weather is to climate, so is the short term to the long term, if we think of forming an outlook or forecast. For example, one theory of climate may explain why regions near the poles tend to be colder than regions near the equator, and this is entirely distinct from a forecast of rain next week in Nairobi or of snow next week at Prudhoe Bay. We may or may not be able to make weather forecasts reliably, but even if we can’t, that doesn’t vitiate our theory of climate that explains the overall temperature differentials across the globe.³ (Then again, our climate theory will require refinement when we discover glaciers on Kilimanjaro, which is almost on the equator.)

³ It is beside the point that in my example, climate change is over space and not over time, whereas weather change is over time. I chose the example of climate change over space rather than over time in order to emphasize the distinction between climate and weather, and implicitly, to suggest that the reliability of forecasts of one is unrelated to the reliability of forecasts of the other.

Similarly, some analysts develop theories of what the stock market may do over the next week, based on valuations, past movements, relationships with bonds, volatility, measurements of investor sentiment, and so forth, but unless you're very, very gullible, you'll deeply discount any short-term market forecast based on such a theory, however plausible or even reasonable. With rather more justifiable confidence (though without arriving at a value precise to the first decimal place), we can attempt a long-term forecast for the stock market. For one thing, we are pretty confident that the long-term return will be positive. How long is the long term? It's not infinite, because, long before the sun cools, swells, boils off the oceans, and expands beyond the earth's orbit, returns will reliably drop 100% and then stay at 0% per year. The long term is not one year, either. But perhaps 100 years is a good long-term span, or 50 years, or even 25 years. If you're a young or middle-aged investor in good health, it might make sense to think of the rest of your life as the long term; if you're an aged investor planning to pass on your estate, then perhaps 10 to 15 years is a good figure for the long term.

Part 2: The Answer

There are two broad categories of methods for coming up with a long-term expected return for the stock market: one is to look to history and to assume that the future will be like the past, and the other is to conduct an economic analysis. Under the latter heading, there are several methods and variations on those methods.

Looking Backward: Projecting the Past onto the Future

If we begin by looking at the history of stock market returns, we'll grow accustomed to the numbers and develop a feel for the possible and the credible.

Most treatments of the stock market's returns begin in January of 1926, for no better reason than our best database of all stock prices and dividends arbitrarily contains nothing earlier.⁴ From the beginning of 1926 to the end of 2010, the annual rate of return on the stock market was 9.9%. If we look back only fifty years, to the beginning of 1961, the annual rate of return was 9.8%, which is very similar. But we do, indeed, have records—not as good, but the result of continually improving data research and management—going back into the nineteenth century, and if we start in 1900, then, over the last 111 years, the annual rate of return was 9.4%. In each case, we're taking 31 December 2010 as the end point. I can remember, ten years ago, looking back to the beginning of 1926 and seeing an annual rate of return of 11.3%. This demonstrates that history, and the long-term record, is not what it used to be.

The most widely quoted expert on historical stock market returns is Jeremy Siegel, of the Wharton School of Business, who latterly has made a career of stock market forecasts based on history. He observes that extended periods of below-average market returns have been followed by extended

⁴ The data were assembled by the Center for Research into Security Prices (CRSP) and were the foundation for *Ibbotson S&P 500 1926-2010 Classic Yearbook: Market Results for Stocks, Bonds, Bills, and Inflation, 1926-2010* (Chicago: Morningstar, Inc., 2011).

periods of above-average returns (and vice versa), and, given that we're emerging from a more-than-ten-year period of below-average returns, he argues that we have good reason to expect that the next five to ten years will see above-average returns. That may be right, but it is an intermediate-term forecast, and our perspective is long term. Professor Siegel's average return, that is, the average that he thinks may be exceeded in the intermediate term, is still the historical long-term rate of return.⁵ (This prediction of a return reversal, based on patterns, presumes that the stock market as a whole is not priced efficiently, in the sense we considered in an earlier essay.)

Given the numbers I just presented, it is fair to say that an annual rate of return between 9% and 10% is a pretty good summary of the historical long term, and, in the absence of any further information, this would be a good estimate for the long-term future. But we do have further information.

Looking Forward: Economic Analysis

In an earlier essay, I explained that the value of an investment should equal the sum of all future cash flows from that investment, converted into the common currency of the present day.⁶ The conversion rates, or, as I called them, the "discount factors," are, in effect, the rates of return. Just as the rate of return converts the present value into the future value, so, too, can the rate of return work in reverse and convert the future value back into the present value. This is true for the stock market as a whole, just as much as it is of any of the market's constituent company stocks.

Valuing the Market's Future Cashflows

Consequently, one way of estimating the future stock market return is to estimate the future cash flows to all the stocks in the market, and to work out what constant rate of return would equate those cash flows to the present value of the market, namely, its price. What is the price of the market? Well, we have indices of the market, like the S&P 500. If we can describe future cash flows not in total dollars, but in dollars scaled to the prices of those indices, we can work out an estimate of the rate of return. Fortunately, there are various commercial services, including Standard & Poor's (S&P), that provide these scaled values of past (and for one or two years in the future) cash flows.

If you have owned a stock, you probably never thought of receiving a "cashflow." You thought you were receiving dividends. So, for a first pass at an estimate of the rate of return, let's try to equate future dividends to the current price of the market. (I'm using the value of the S&P 500 on December 31, 2010, which was 1257.64.) The return that I find is roughly 6.2%. That's very low.

But future dividends don't represent all the cashflows that an investor might receive from a company or from the market. Sometimes, companies return value to shareholders by buying up their stock in the market. Sometimes, companies hoard cash for a long time, before finally relenting

⁵ See the summary of Jeremy Siegel's current thinking in Gene Epstein, "Equities Investors, Don't Despair," *Barron's*, March 21, 2011, pp. 36-37. Siegel is famous for his book *Stocks for the Long Run*, 4th edition (New York: McGraw-Hill, 2008)

⁶ Peabody River *Newsletter*, issue 7, April 2010, *op.cit.*

and returning it to the shareholders as regular dividends, a one-time dividend, or a share repurchase. Warren Buffett's Berkshire Hathaway famously pays no dividend at all. It's really a matter of what a company *could* pay out in dividends, not what it actually pays out. So, at the upper limit of what the future cashflows might be, we can redo our calculation on the assumption that the value of the market equates the value of all future earnings, of which only some portion may be paid out in dividends, to the current price. The return that I find when I do this is roughly 10.8%. That's in line with the long-term historical return of the market, though a bit high.⁷

But we can be quite sure that not all of a company's (or the market's) earnings will ever be returned to the shareholders, who have claim to only a portion of the earnings, so our 10.8% estimate of the market's return is an upper limit. And the 6.2% return estimated from dividends is a lower limit. Let's split the difference and shade it down, and say that (according to this particular methodology), our estimate of the future return on the U.S. stock market is about 8%.

That range of 4.6 percentage points (between 6.2% and 10.8%) is very wide and could result in a stupendous difference in amounts of wealth at some point in the future. It shows up the high degree of uncertainty in our forecast. Furthermore, I spared you the details of my calculations, but you can be sure that my estimates of future dividends and future earnings (or the rate of growth of dividends and earnings from current levels) were very rough indeed. Still, there is a lesson here beyond the uncertainty: *The range of estimates is mostly below the long-term historical rate of return on the market.*

But this isn't the only way of estimating the future return on the stock market.

Introducing the Equity Risk Premium

One alternative way of thinking about the stock market's rate of return is that it is a layer cake of returns, one stacked atop another. In its simplest form, this idea suggests that the market's rate of return has two layers: the rate you would earn from a riskless investment, and the rate the market pays you for taking on the risk that things might not turn out well. This second component is called the **equity risk premium**. (Remember that when I say that the market is paying a return, I don't mean that the market has a man behind a curtain setting prices; rather, the net result of institutions and individuals, like you and me, judging prices and trading among themselves in the market is that stock prices have to compensate these participants for taking risks.)

There is, indeed, a riskless investment, which has a more-or-less knowable rate of return. U.S. government Treasury bills, which pay you back after three months, are as close to absolutely safe as it is possible to find in this corruptible world. (Forget the silly political persiflage you may have heard of the U.S. government's being bankrupt.) So the rate of return on Treasury bills is what financial economists usually regard as the riskless rate of return. But because Treasury bills are good for only three months (after which the money is returned to you and, presumably, you reinvest it in new Treasury bills), there's a mismatch to our time scale for long-term market returns. Instead, for our purpose of estimating the long-term stock market return, it may be more fitting to take as our riskless foundation the U.S. government twenty-year bond (which matures in twenty years), which is

⁷ The methodology that I am using in these calculations is my own, and it is beyond the scope of this essay to present it. I will provide details upon request.

equally free of the risk of default. And it is generally the case, during most short periods of time, and always the case over long periods, that the longer a bond takes to mature, the higher the interest rate on it.⁸

So, we can use historical data to estimate the equity risk premium as the difference in rate of return between the stock market and the U.S. government twenty-year bond. And if we take the period from 1925 to the end of 2010, we find that this annual difference averages 4.0%.

One especially useful feature of bonds, which is not shared by stocks, is that (assuming that the issuer of the bond doesn't default), you know what rate of return you will have realized when the bond reaches its maturity date.⁹ So, I can look online and see that right now, the twenty-year U.S. Treasury bond has an interest rate of about 4.8 %. I add to this the equity risk premium, and—Presto!—I have a forecast of 8.8% for the stock market's rate of return over the next twenty years.

But projecting the long-run historical equity risk premium into the future is nearly as crude and uninformative as projecting the historical return. The equity risk premium has varied considerably over time. That's the whole point of an equity risk premium. If it were certain, there would be no risk, and its value would be 0. So our original question could, with scarcely any change of meaning, be reworded as: "What equity risk premium can we expect over the long term?"

Our little calculation has served mainly as an introduction to the concept of an equity risk premium.

Forecasting the Equity Risk Premium

So the best way to forecast the market's rate of return is to apply economic arguments to estimate the equity risk premium, and then to add this to the risk-free rate of return. And the best version of this method that I have seen depends explicitly on the limit to the economy's growth.

It stands to reason that the companies whose stocks are traded in the stock market cannot forever grow faster than the economy, because, if they could, all the companies that constitute the market would be bigger than the economy that contains them. If you pay attention to the nightly business reports on the radio and television, you receive frequent updates on what the growth rate of the U.S. economy has been, and what the growth rates of other economies, especially China's, have been. If you pay more attention to the glamorous Chinese economy than to the U.S., you may forget what is a realistic long-term growth rate for a mature economy, like ours. It's probably in the vicinity of 3% (net of inflation) a year, not the 8% to 10% we've been seeing in China. (The White House and the Congressional Budget Office are currently forecasting slightly less than 3%.) Using the assumption of 3% per year for the U.S., Bradford Cornell, an economist at CalTech, has worked out that the long-term equity risk premium should be no more than 4% to 5% per year, after allowing for

⁸ I am simplifying. Most economists would say that there is the risk-free rate of return, on top of which is a return premium to compensate for the longer time horizon.

⁹ I am simplifying again. Strictly speaking, this is true only on the assumption that the dividends that the bond pays are reinvested at the same return that the bond itself is earning, which is unrealistic. It is, however, good enough for a quick estimate. The return calculated on this assumption is the yield-to-maturity, and is the usual way of quoting the value of a bond.

inflation.¹⁰ Let's add this to the twenty-year risk-free rate of return, which as we saw is currently 4.8%. *Then Cornell's estimate is in the range of an upper limit of 8.5% to 9.8% per year.*¹¹ He specifically states that this is an estimate of the upper limit, not the expected value. Notice that Cornell's value appears to be in line with the historical value. But that was an actual value, and Cornell's is an upper limit.

Cornell's paper was published in early 2010, but his calculations were as of the end of 2008, when the stock market was much lower than it was at the end of 2010, which is my reference point.

If the market were **efficient**, meaning that it always priced stocks correctly, the market could not, by definition, be priced too low with respect to its true value. But as I explained in an earlier essay on market efficiency,¹² there is reason to suspect that, while the market is very good at pricing stocks relative to each other, it may sometimes get their value in aggregate wrong. Cornell believed that the market was underpriced at its recent depth, close to the time he was writing, and that it would therefore bounce back. This future bounce is reckoned into his calculation. But since the market bounced a cumulative 46% over 2009 and 2010 (and more in the last few months), some of his equity risk premium has already been realized. If we were to repeat Cornell's calculations today, our upper limit on the forecast rate of return for the market would be lower.

So, our economic analyses have given us estimates of the stock market's long-term future return of between 6.2% and 10.8% (and probably less than the middle of that range, around 8%); of about 8.8% (by a method that is mostly a projection of history); and with an *upper limit* of 8.5% to 9.8%, based on the most rigorous analysis, but with outdated information, which, if updated, would almost certainly result in lower numbers.

In light of these values, I feel most comfortable with a forecast return of about 7% for the long-run return on the stock market, though it may reflect a soupçon of pessimism.¹³

So much for return; what about risk?

The vagueness of our estimate of expected return doesn't fully comprehend the risk that things might turn out worse. There's a subtle and potentially confusing distinction here. We need estimates of two quantities: the expected rate of return, which is our best estimate of the future rate of return, and the measure of expected risk, representing what might be wrong with our estimate, and how bad the stock market return might turn out to be.

¹⁰ Bradford Cornell, "[Economic Growth and Equity Investing](#)," *Financial Analysts Journal*, Vol. 66, no. 1, January/February 2010, pp. 54-64.

¹¹ Cornell actually adds his equity risk premium to the short-term risk-free rate, not the twenty-year rate.

¹² Peabody River *Newsletter*, issue 8, July 2010, "[Is the Market Efficient?](#)"

¹³ An important paper published in 1985 by Rajnish Mehra and Edward C. Prescott pointed out that the observed historical equity risk premium was far larger (by an order of magnitude) than financial theory says is justified by the risk of stocks. (Mehra and Prescott, "[The Equity Premium in Retrospect](#)," NBER working paper 9525, National Bureau of Economic Research, February 2003.) The very idea that there might be an "equity risk premium paradox" was so bizarre (for economists, who were the only people who knew of the paradox and cared), that Mehra and Prescott had some difficulty getting their paper published in a proper economics journal, though they eventually succeeded. It is now accepted that the paradox is real, but no attempt to resolve it has been entirely successful. I'm not implying that because the theory must be right, our estimate of the equity risk premium is too high. But it still gives me pause, and it wipes away the fear that I am being unduly pessimistic.

Recall our initial definition of the expected rate of return, which related it to the myriad possible ways the annual returns might be realized. We've just arrived at our expected rate of return without having to calculate all the infinite possible sequences of annual returns to the future, along with their probabilities. Conceptually, however, it still fits our definition of expected return: Implicitly, if the multitudinous possible market trajectories were to be weighed in the balance, our rate of 7% per year represents the return around which the greater possible rates of return are in equipoise with the lower possible rates of return.

When we look at the expected return this way, we see it situated in the context of investment risk, the possibility that the return (and our investments) will turn out badly.

If you were ever a student in a first-semester statistics class, you may recall that your teacher or your textbook would bestow on you, for problem sets, the expected value and the amount of variability as two precise numbers, as if they were knowable realities. But we're in the dark cave of investing, and we see only the shadows of these values on the cave wall.¹⁴ The vagueness of our estimate of the expected rate of return reflects its unknowability, which is not the full measure of the investment risk. The quantity of risk itself is also, ultimately, unknowable.

I know of no better way of estimating long-term future risk than to look backward at the variability of past returns, though there are complementary ways of interpreting the historical numbers. There are no economic methodologies for estimating long-term future risk. (There are economic and statistical methods for estimating risk over the short-term future, but our time scale is vastly greater than is appropriate for these methods.) We know that we can lose a lot of money in the stock market in one year; in 2008 alone, the stock market lost -37%. But year-to-year fluctuations tend to cancel each other. (Not entirely though; it's a common error to think so, and I'll have more to say on this in a moment.)

Let's consider a **holding period** of ten years. That is, we assume that we invested some money in the stock market and held onto that money for ten years. We'll start by looking backward from the end of December 1935, because we have values of actual returns going back to the beginning of 1926. We find that over those ten years, we earned an annual rate of return that was 3.8% better than if we had just kept our money in cash (or, more precisely, Treasury bills, which, as distinct from loose change, is what an economist means by "cash"). We're once again looking at the actual equity risk premium. Now we move forward to the end of January 1936 and look back ten years again. We keep stepping forward a month at a time, and looking back ten years. This way, we have 901 **rolling ten-year holding periods**. Of these, 150 had stock market returns that were worse than just staying in cash. We can repeat this exercise for rolling twenty-year holding periods. There were 781 such periods from 1926 through 2010, and during these, the stock market always came out better than being in cash. But sometimes not by much: For all that stock market variability, there were 27 twenty-year periods where the annual rate of return on the stock market beat cash by less than a

¹⁴ In more technical but clearer language: To define the uncertainty of future returns, we're trying to establish their statistical distribution, which has at least two parameters (expected value and standard deviation), and almost certainly more. But not only do we not have the population values of the parameters; we don't even have decent samples from which we can estimate sample values of the parameters. And we don't know what distribution we're describing, though we're confident it resembles a lognormal distribution, at least away from the tails.

percentage point. Less than a percentage point per year can add up to a significantly different amount of money over twenty years, but still, this isn't a spectacular outcome.

You should not infer from this history, however, that if your time horizon is twenty years or more, you can be certain of coming out ahead if you put your money in the stock market. The market hasn't signed a notarized contract to deliver to you a positive return over any time horizon, and the only natural constraint on its bad behavior is that can't lose more than 100%.

An alternative way of getting at the long-term risk of the stock market, still using history, is to infer the chances of specific bad outcomes from the frequency with which various returns occurred in the past. That is, given the market's past behavioral patterns, we can infer the likelihood of its misbehavior in the future.

If the pattern of past monthly or annual returns looks like a bell curve (to take one very convenient pattern), then well-established mathematical properties of the bell curve will allow us to calculate the probability of an annual rate of return, for any chosen time horizon, that is worse than any arbitrary cut-off we might establish.¹⁵




Figure 1

What, for example, is the probability that the annual rate of return will be less than zero over a twenty-year time horizon? To answer this (on the assumption of a bell curve), I need the two numbers my statistics teacher would give me: the expected annual rate of return above the return on cash, and a value representing the volatility. I shan't go much into the latter,¹⁶ which is too technical

¹⁵ This is much, much too crude, because we know that there are upper and lower limits. If I were striving for accuracy, I'd use a lognormal curve. But the lognormal curve has its own faults, making a return of -100% mathematically impossible though it is economically possible, and having an infinite right tail, which is economically impossible.

¹⁶ For the mathematicians: It's the standard deviation of the annual return premium of stocks over cash from 1926 through 2010, converted to a continuously compounded value.

for this essay, but for the expected rate, let's use our estimate of 7%. Subtracting the historical average rate of return on cash, we get 3.3%. Then, using the standard mathematics of the bell curve, I calculate that there is a probability of 21.5%, which is to say, about a 1-in-5 chance, that over the next twenty years, an investment in stocks will turn out to be worse than just keeping money in cash. Because the risk changes with time horizon, the corresponding figures for ten years are 28.8%, or a little worse than a 1-in-4 chance, and for thirty years are 14.4%, or about a 1-in-6 chance. So, even though we've never seen a twenty-year span, let alone a thirty-year span, in which the stock market's rate of return was worse than keeping money in cash, we can calculate the probability of such a risk in the future.

If you have been paying close attention, you'll have noticed that the 1-in-4 chance for a ten year interval is much worse than what we saw in the past, 150 cases out of 901 (which is only about 1-in-6). Is all our complex mathematics overestimating the risks? Yes and no.

In an earlier essay, I explained that the familiar bell curve is not a good model for the behavior of stock returns.¹⁷ In particular, it leads to overestimates of the small risks, but underestimates of big risks. A rate of return that matches that of cash may not be so terrible. So, it is, indeed, possible that our 1-in-6 estimate of the chance of a stock market rate of return worse than cash for a thirty-year span is an exaggeration of the risk. But if some possibilities are less probable, then other possibilities have to be more probable, and it is correspondingly likely that the chance of a really bad outcome, like having a result worse than cash by *at least* -5% per year (which means, relative to cash, a cumulative decline of -79% over 30 years!), is greater than I'd calculate it to be.

Table 1 presents the risks of bad outcomes relative to keeping money safe in cash, for different time horizons, and different standards of what a bad outcome might be:

		Return Relative to Cash or Worse, per Year		
		0%	-3%	-5%
Horizon, in Years	10	1-in-4	1-in-7	1-in-13
	20	1-in-5	1-in-16	1-in-50
	30	1-in-6	1-in-33	1-in-170

Table 1

When contemplating this table, bear in mind that these values depend on our estimate of 7% for the expected rate of return on the stock market (and, for that matter, on the historical rate of return of 3.6% on cash), and on historical statistics for volatility. A higher expected return on the stock market would result in less risk, a lower expected return in higher risk. And because they assume that the pattern of possible future returns resembles a bell curve, which we know is incorrect, the values in the lower right of the table are probably too low; that is, the risks are greater than they suggest. Also, remember that the returns are rates; -3% per year over ten years corresponds to a *cumulative* decline of -26.3%, and -5% per year over ten years corresponds to a *cumulative* decline of -40%.

¹⁷ Peabody River *Newsletter*, issue 3, January 2009, "[How to Think about Investment Risk.](#)"

Monte Carlo Simulation

There's yet another statistical method of estimating the risks of the stock market. Although I won't squeeze a result out of it now, it's worth knowing what the method is. It simulates the stock market's possible returns over time, and starting about ten or twelve years ago, it has been heavily used by financial planners. It is the computerized equivalent of the following: write each historical annual difference between the stock market's return and the return on cash on a slip of paper,¹⁸ put all the slips into an upturned hat, and stir them around. Take out slip, write down the number as the return for year 1, put the slip back into the hat, and stir. Take out a slip, write down the number as the return for year 2, put the slip back, stir. Repeat this for each year through year 30. You now have a possible sequence of returns for the next thirty years. Calculate the rate of return that corresponds to this sequence. Repeat this exercise a few thousand times, so that you have a few thousand possible sequences of returns for the coming thirty years, and a few thousand possible rates of return. This exercise is called a **Monte Carlo simulation**. If you sort through the resulting rates of return, they'll form a pattern, something like the bell curve, but not exactly, because the pattern will reflect the way the actual historical data deviate from the bell curve.¹⁹

With this curve, you will be able to count and see what percentage of all the rates of return were below 0%, below -3%, and so forth, and the resulting risk estimates might be fairer than those of our purely analytical computations.

Expected Return and Risk Behave Differently over Time

Notice that long-term expected risk behaves differently from the long-term expected rate of return. The expected rate of return over the short term should be the same as the expected rate of return in the long term, at least in the absence of information to the contrary. But also in the absence of additional information, risk, if defined as the likelihood of a rate of return below a given benchmark, decreases over time. The short-term outlook for risk (by our definition) is greater than the long-term outlook.

Notice also that all our methods of estimating risk, even the count of historical twenty-year periods, treat risk as purely a game of numbers, without regard to causes. But, as we assumed when we calculated the expected rate of return on the stock market, returns are economic phenomena, created by expectations of cashflows and suitable discount factors. If the market behaves badly, it will likely be the reflection of a broad outlook among traders for dire problems with the economy at large. Seen this way, the statistical analysis of risk takes into account only symptoms, not underlying causes, and if you firmly believe that the economy has next to no chance of facing a disaster, then our estimate of a 1-in-170 risk of a cumulative decline of -79% with respect to cash over the next thirty years will, in your judgment, be much too high.

¹⁸ To adjust the results of this exercise to reflect our forecast rate of return on the stock market, you should also subtract from every number the difference between the historical rate and our forecast rate.

¹⁹ There's a subtle complication at work here. The simulation requires that I specify an expected return for stocks, but this will actually be the expected *median* return of the multitude of simulated possibilities, which is not the same as the expected return. The median is the value that is in the middle of a range, with half the values falling on either side. For stocks, the median value is never the same as the average that is the expected return.

All the estimates in this essay are based on either economic or statistical reasoning. If you are in thrall to the occult and believe that the stock market is driven by forces not based in economics, like astrological harmonies or mysterious patterns of long waves, you may come up with very different estimates of return and risk, but I daresay that ultimately, you'll have to agree that the growth of the market will be limited by actual economic forces that constrain the size of the companies it represents. And, of course, on the downside, returns are constrained by a cumulative value of -100%. Unless you choose to go into debt, you can't lose more than everything with which you began.

But Wait! There's Less!

The stock market returns that we have been estimating are returns *before* taxes are paid on them. If you invest in the stock market through a taxable account, your returns will be reduced by income taxes. There are all sorts of clever and legal ways of avoiding or reducing taxes on investments, such as balancing realized capital losses against realized capital gains, but our interest here has been in ideal theoretical outcomes for the market, not the results of clever investment and tax strategies. A straightforward investment in the stock market through a taxable account will certainly have its returns reduced by, at the very least, the income tax on qualified dividends, currently 15% at the Federal level. Moreover, income taxes can add insult to injury. In 2008, when the stock market produced a return of -37%, it nonetheless paid dividends. Dividends always represent a positive return, so, if we ignore those, the *price* return in 2008 was worse, -38.5%. But the dividends, which slightly counteracted the market decline, were taxed. After applying a tax rate of 15% to the dividends received that year, the after-tax total return on the market in 2008 was actually -37.3%. If our expectation for the long term is of a rate of return of 7% on the stock market, and of that return, 2% comes from dividends, then the after-tax return on the stock market will be about 6.7%. That reflects just Federal tax; your state may impose an additional tax on the dividends. And it doesn't matter if you, as an individual (as distinct from my readers who represent non-profit, and therefore tax-exempt, organizations) have invested your money through a tax-advantaged account, like an IRA. When you withdraw your retirement money, you'll be paying tax on it.²⁰

The results are diminished further when you recall that all the values of the market's rate of return that we are considering include inflation. Take out inflation to reflect the actual purchasing power of the returns, and you find that the market's expected rate of real return is actually much smaller. So if inflation is, say, 2.5% per year, then the real expected after-tax return on the market is down to 4.2%.

The Typical Return

The expected annual rate of return isn't what you might call the "typical" annual return. The typical return that you might expect is the average of the returns that you'll see over the coming years. As I noted near the beginning of this essay, the average is greater than the rate (unless there's no variation

²⁰ At the level of "mental accounting," this is a little ambiguous. Without the tax advantage of a traditional IRA, you would have paid tax on the earned income that you stashed away before you stashed it away, so the IRA is allowing you to defer this tax, and is not a tax on the investment growth as such. But alternatively, you need to be aware, as you see your IRA grow, that not all its value will be coming back to you.

from year to year). A simple mathematical method allows us to estimate the average, or typical, return of the market, given the rate of return and the volatility of returns. Using this method, we can convert an annual rate of return of 7% into a typical annual return of about 9%.²¹ That's nearly 3 percentage points less than the historical average annual return of 11.9% since 1926. In 2010, the market returned even more, 15.1%. Of course, we may seldom observe the precise typical return, but, given the nature of the bell curve, actual one-year returns that are close to it should occur with some frequency.

Summary

To recapitulate the forecasts in one place: We find that the expected annual total rate of return on the market is probably around 7%, though our uncertainty of this estimate is such that we've seen a range from 6.2% to 10.8%, neither of which is very credible by itself. The expected long-term rate of return is less than the historical long-term rate of return, which, depending upon our choice of historical period, has been between 9% and 10% per year. Furthermore, Federal taxes, to say nothing of state taxes, will lower the expected rate for stock market investments held in taxable accounts; if the forecast is 7%, then the after-tax forecast is about 6.7% or less. If we take out an assumed inflation rate of 2.5%, then the pre-tax rate is about 4.5% and the after-tax rate is about 4.2%. The risks may be represented very crudely by Table 1, though the numbers there probably underrepresent the true risks of disastrous outcomes. And the typical annual return that we can expect to see in the future is about 9%.

For the investor, stock market return is not destiny. On the one hand, there are investments, like bonds, that we normally expect to have lower returns (thanks to lower risk), but there are also investments that may well have higher returns. The rate of return on the stocks of smaller companies, in aggregate, has exceeded the rate of return on the stock market as a whole, and there is reason to believe that this will continue, though the returns are at the expense of greater risk. It's entirely possible that the returns to the stock markets of emerging economies will also be larger than the returns to the stock markets of developed economies, at least until those economies emerge. And, before they do, there are the so-called "frontier markets," like Pakistan's and Tanzania's, whose returns are even riskier than those of the emerging markets.

Ingenious stock selection, if someone can succeed at it, or clever strategies that actually work in practice and not merely on paper, or even just a combination of risk-taking and good luck might produce much higher returns than the market. For example, if you had had perfect one-month foresight, so that you sold all your stocks immediately before months when the stock market went down, and then bought back into the market immediately before months when it went up, your stock market investment from 1926 through 2010 would have earned an annual rate of return of 32.7%.

²¹ For the mathematicians: The expected value of the natural logarithm of $1 +$ the growth rate can be expressed as a Taylor series of the natural log of $1 +$ the expected average return plus terms representing successively higher statistical moments, also in terms of the expected average return. For well-behaved statistical distributions of returns and small values of return, the series approximates to the expected average return minus half the variance. See Jarrod Wilcox, "Harry Markowitz and the Discretionary Wealth Hypothesis," *Journal of Portfolio Management*, Spring 2003, p. 61. I used the long-term historical variance of stock returns, about 0.04, in my calculation.

Regardless, however, of whether you think you can get greater returns through straightforward proven methods, or through unproven and improbable ingenuity, the long-term rate of return on the stock market is a benchmark against which we will measure our expectations, and just maybe, around which we will plan our economic futures.

Adam Jared Apt, CFA

Disclaimer: Nothing in this newsletter should be construed as specific investment advice or a recommendation to buy or sell securities. Any investment advice given here is general in nature and offered without expectation of compensation.