



How to Think About Investment Risk

By Adam Jared Apt*

January 27, 2009

Advisor Perspectives welcomes guest contributions. The views presented here do not necessarily represent those of Advisor Perspectives.

This introductory article is intended for the educated layman. It bypasses consideration of certain kinds of risk, such as liquidity risk, default risk, and systematic risk.

If hopes were dupes, fears may be liars.
Clough

We are in a dark wood, and the way is not straight. The way of investing never is. But try to put out of your mind, for the next few minutes, the present dismaying state of the economy and your personal finances. My subject is the meaning of risk in the context of investing. All concepts and ideas have a history of genesis and development, and just so do our ideas concerning risk; the ideas themselves ought nonetheless to be timeless. The concept of risk, as an abstraction, should not depend upon the level or direction of the stock market.

Let's first consider a risk that isn't related to investing. If you have read Bill Bryson's 1998 book, *A Walk in the Woods*, his account of hiking the Appalachian Trail, you will remember that at the outset, he dreads encounters with bears. At one point, he reassures himself with the historical fact that no one has been killed by a bear in New Hampshire or Vermont since 1784. A few pages later, though, his confidence evaporates, and he asks, "And how foolish must one be to be reassured by the information that no bear has killed a human in Vermont or New Hampshire in 200 years? That's not because the bears have signed a treaty, you know." Please pardon the metaphor, but that last remark is an important theme of this essay. Where there is uncertainty, there is risk. And there is uncertainty in finance because the bears have not signed a treaty.¹

Risk as the Possibility of Loss

Long before the rise of modern financial theory, investors were well aware that their investments were at risk of losing value. In the middle of the last century,

¹ As of this writing, the statement about Vermont and New Hampshire is still true. In 2000, however, a hiker in the Great Smoky Mountains was killed by a mother black bear and her cub.



Benjamin Graham, the father of securities analysis, put forth the concept of “margin of safety,” which he described as “the secret of sound investment.” I will oversimplify his definition, but for a stock, it represents the earning power of the company that issued the stock, regardless of the stock’s price. That is to say, the margin of safety is a measure of reassurance that, should the stock’s price fall very low, it would continue to be worth something because of its stream of future earnings. According to Graham, you want to purchase investments with large margins of safety. But though he quantified the margin of safety, he did not quantify risk, which he defined as the possibility of the loss of value realized at the time of sale of an asset.

With the rise of modern financial theory, financial risk has come to have many definitions, not necessarily because analysts have philosophical disagreements—though they do, and I will mention them—but because we now recognize many kinds of risk. Nearly all are or can be expressed in the form of a probability or chance of loss. The risk that is most pertinent to an individual investor, I believe, is the possibility of not having the expected amount of money when that money is needed, a definition that is close to Graham’s. Retirement accounts offer an especially stark illustration: If a couple requires and expects, say, a nest egg of \$3,000,000 for retirement at age 67 in order to live out the remainder of their lives in the manner they choose, then they are at risk to the extent that there is a possibility that they will have less than that amount. But there is also risk in the case of the couple who wish to pay for a private school for their child out of their savings. One should not impertinently dismiss the latter as a luxury while considering the former a necessity. The latter couple may have concluded that their child was more likely to flourish in the private school than in a public one, and the former couple would be aiming to live at a level of comfort that is unknown to many poorer retirees. And perhaps it is the same couple that has both goals; that is, a single person or family may face multiple risks and have to face tradeoffs. Both plans are subject to risk. An investment advisor can help clients work through the problem of addressing these goals.

But this definition of risk, though reasonable, is not very practical by itself. How would I estimate the chances of having insufficient funds to pay for these requirements? And what is the relevance of this definition when we talk about the risk of a single stock, or the risk of the stock market? My holding of, say, Ford stock or Citigroup stock does not all by itself carry the risk that I won’t at some point in the future have the funds to pay for my retirement. That’s because I hold other investments as well, which may be able to pay for my plans.


Risk as Volatility

Economists and investment analysts therefore use alternative definitions for describing the risk of a particular investment, such as a stock, bond, or mutual



fund. The most basic and common, and the one that is taught in introductory finance courses, is that risk equals volatility of returns. Volatility can be treated mathematically. If you can measure it, this mathematical tractability allows its use for prediction of future uncertainties. Moreover, this definition has an intuitive appeal, because it is undeniable that many investors alternate between elation and despair as the market (or particular stocks) go up and down.²

Once you define financial risk as volatility, well-understood tools of probability and statistics permit you to plug a few values (parameters) into established mathematical formulæ and to calculate the chances of a rise or fall in the price of a stock, or the market, or a bond, or your entire portfolio, and therefore your chances of not having enough money when you need it.



Here’s how it works. On the left, you see a plot of the distribution of annual returns to the U.S. stock market (as represented by the S&P 500 index) between 1926 and 2008. It’s a graph of the number of times that returns fell within defined intervals. You see, for example, that there were very few years with returns that were less than -30% and only a couple of years with returns that were more than +50%, but there were a lot of years with returns that were between 10% and 20%; the average is 11.67%.³ The shape of this graph looks vaguely like the graph on the right, the familiar “bell curve,” which is also known as the “normal distribution.” The bell curve is described by a mathematical formula. This formula requires that two, and only two, values be plugged in for the entire scale and shape to be defined. One of these values represents the middle or peak of the curve (the average) and the other represents how spread-out the curve is. For

² Note that I am referring to the volatility of returns, not of prices. There are several reasons for preferring return volatility to price volatility. The most obvious is that returns are percentages, and therefore can be compared across different investments. The price volatility of a stock whose price averages \$6 per share is not comparable to the price volatility of a stock whose price averages \$100 per share.

³ Data, except for 2008, from *Ibbotson S&P 500 Classic Yearbook: Market Results for Stocks, Bonds, Bills, and Inflation 1926-2007* (Chicago: Morningstar, Inc., 2008).



investment returns, the “spread-outness,” so to speak, is the volatility. (If volatility were low, most returns wouldn’t vary much, and they’d be tightly clustered around the average, and the curve would be very tall and narrow; if volatility were very high, the curve would be very spread out, with a low peak.)

So, for example, given the actual historical returns that underlie the graph on the left, I can calculate the average and the “spread-outness” and plug them into the formula, which then tells me that there is a 1-in-20 chance that the return on the U.S. stock market will be less than -22.2% in a single year. (I can calculate the return corresponding to any chance I choose: 1 in 10, 1 in 50, 1 in 1000, and so on.) With a just a little more mathematics, this can be pushed a bit further to show that if you invest \$1,000,000 in the U.S. stock market today, your investment stands a 1-in-20 chance of being worth less than \$1,101,873 in ten years. In a wry sort of way, this is comforting information; we can reword that last statement to say that we ought to expect that in only one decade during two centuries, your portfolio will grow by less than 10.2% in value. A total of 10.2% growth over ten years corresponds to an average rate of growth of 0.97% per year.⁴ That’s not enough to compensate for expected inflation, but at least it’s positive.

We’re getting a lot closer to being able to deal with the real risk to your financial wellbeing.

Equating investment risk with volatility has been extraordinarily fruitful. Much of the structure of modern finance rests upon this definition and its implications. And by “finance,” I mean not just the theory and practice of investing, but also corporate finance, the way companies finance themselves, evaluate decisions to pay for the creation of new products and factories, and buy and sell businesses. It also applies to the ways that financial institutions, primarily banks and insurance companies, estimate the business risks that they will face in the event of a credit crisis.

Recall Graham’s “margin of safety.” Modern financial practitioners, using the mathematical tools of probability and statistics, have developed a different form of margin of safety, called “Value at Risk,” or VaR (pronounced “varr”). Much as I just calculated that there was a 1-in-20 chance of a \$1,000,000 portfolio being worth less than \$1,101,873 in ten years, the risk managers—mathematicians who are employed by banks and other financial institutions for just this purpose—can calculate how much a bank’s capital or an insurance company’s reserves, say, might drop in a month’s time with a 1-in-1000 probability (meaning that the bank might experience such a loss in one month out of eighty-three years), or a 1-in-10,000 probability. The calculations are much more complex than the one

⁴ This is the compound rate of growth, which is smaller than the arithmetic average.



underlying my example, but the basic concept is the same. Recognizing that the bears haven't signed a treaty, and that there is always some chance, however small, of a catastrophic loss, they can now determine what that loss (or worse) might be for any probability they select. Whereas Graham and his followers recognize that some risk exists, and concentrate on ascertaining that there is a sufficiently large cushion to absorb the impact of a big loss, modern financial practice (for banks and insurance companies) explicitly calculates the size of a loss that can occur at any selected probability, and adjusts the portfolio to ensure that the worst losses are very unlikely.

Simple Misunderstandings

Before I reveal how this kind of reasoning has led investment managers down the primrose path, let me respond to several common objections to the definition of risk as volatility.

Some financial novices think they have found a "gotcha" when they observe that volatility that produces returns greater than the average isn't risk. But that's not a problem. Forget the upward moves; as long as we can calculate the probability of downward moves, we're still dealing with risk.

Although I told you that the graph on the left looks like a bell curve, you ought to have noticed that it doesn't look exactly like one. For one thing, it is slightly asymmetrical. I can tell you that if we were looking at monthly rather than annual returns, the graph would look more like a bell curve, but the "tails" of the graph would be fatter than those of a true bell curve.⁵ Fat tails mean that very large and very small returns occur more frequently than the bell curve formula implies they should.

Again, this is not a serious difficulty, at least for approximations. Considerable mathematical expertise has been devoted to working out exactly what curve best describes investment returns. The familiar bell curve doesn't fit the data as well as some other curves (that is to say, other probability distributions), but that only means that the mathematics required for the calculations is a lot harder, not impossible.

Some investment advisors who have never cottoned to modern finance object to the treatment of risk as a statistical phenomenon, as if the vagaries of the returns of a stock of a specific company or even the returns of the stock market as a whole, behaved in a random fashion, and without cause. But this is to confuse the manifestation of volatility with its causes. Prices may move for a reason (though many of us, through long and bitter experience, become sufficiently

⁵ The shape of the plot changes somewhat for different frequencies of returns: annual, monthly, daily, hourly.



cynical to question this). But the economic and business forces that move stock prices aren't necessarily regular and predictable; hence, price movements, and returns, appear to be random.

Deeper Misunderstandings

Not all misunderstandings of the equation of risk with volatility are the fault of those who object to the definition. Some arise among those who love the definition not wisely but too well, who are enamored of the concept from their first introduction in business school. They have conveyed this misunderstanding in watered down form to the public. It comes out in their bland assurance—I'm sure you've been hearing this constantly over the last few months—that stocks and other investments may go up and down, but if you hold on tight, you'll always come out ahead in the long run. They visualize volatility as just a series of waves around a continuing upward trend. From this, some members of the public cheerfully draw the fallacious inference that their portfolios will have higher returns if they just crank up the riskiness.

But there should be no inference from the definition that the waves of volatility—if they are waves—are regular or easily predictable. They can be small or enormous. It is also a mistake to think that the short-run fluctuations of returns must cancel each other out in the long run. (Graham seems to have made this conceptual error.) They very much tend to do so, and I take comfort from this, as should you, if you can wait a long time. As long as the economy continues to grow, and the government doesn't nationalize businesses, and businesses don't broadly replace common stock with other forms of financing, the stock market should trend upward. But there is neither economic law nor mathematical proof that stocks (or other investments) must always do so. There is always the possibility that, over time, downward tumbles will pitch us into a slough, even with repeated efforts to climb back up, and the trend of returns will turn out to be very much worse than the expectation. In other words, the average return that we experience may turn out to be very different from the average we expected over any future span of time that we may specify. By analogy, bears have been known to stalk their victims; they don't always appear of a sudden. Even if the long-term trend of stock values is ever upward, that is small consolation to those who need to pull out their money when valuations are in a trough.

And who has been peddling this nostrum? Well, many financial advisors and even economists who should know better—I've heard them—have done so. A business journalist or other public dispenser of investment wisdom may have reassured you, for example, that there has never been a twenty-year span when the U.S. stock market has not outpaced inflation, which is true. But, to repeat, the bears haven't signed a treaty. As the great economist Paul Samuelson once quipped, we've observed only one long run, and "one" is not a statistical sample.



(You can't make statistical inferences from one observation, and it's the future long-run average of returns and their dispersion, not the past long-run average and dispersion, that concern us.) Actually, we've observed multiple long runs, if you consider foreign stock markets. And that statistical sample, at least of stock markets in the twentieth century, provides less reassurance. There were a number of stock markets that went out of existence altogether, such as the Russian stock market after 1918. Furthermore, that's an entire stock market. You can recall as easily as I the names of individual corporations that have shuffled off this mortal coil. Lehman Brothers was a large company; Enron was much larger.

The Real Problem with Volatility

Did you snicker, several paragraphs back, when I said that banks and insurance companies can calculate the size of a possible loss at any specified probability? You were right to do so.

Among the causes of the current financial crisis was the failure of our financial institutions to calculate correctly the risks inherent in the portfolios of financial securities that they held. This is not so much the fault of identifying risk with volatility, as the result of the difficulty of applying this definition in practice.⁶ This difficulty humbles the manager of your personal portfolio no less than it does the senior executives of the largest banks.

The source of the difficulty is an excessive but unavoidable dependence on historical numbers for predicting the future. Graham, when expounding on the importance of the margin of safety, was sensitive to this problem; that's why he emphasized the importance of a very large margin of safety. "The bond investor does not expect future average earnings to work out the same as in the past," he wrote, "if he were sure of that, the margin demanded might be small."

When, a short while ago, I blithely calculated that there was a 1-in-20 chance that the U.S. stock market would have a return of less than -22.2% in a single year, I

⁶ The recent collapse of our financial markets had many causes, but its scale would have been smaller and the consequences mitigated if the banks, investment banks, and insurance companies had had a firm grasp of the risks that they—and we—confronted. I had thought at first that the risk management departments, being cost centers and not profit centers, were being ignored by senior management eager to make lucrative sales. I still suspect this, but my reading now leads me to think that at least as important were the inability of the risk management departments to get a thorough accounting of the securities and obligations that the firms held, and their reliance on precise VaR estimates that were predicated on fanciful models of risk. On this last point, and for a profound exploration of the concept of investment risk, I strongly recommend Riccardo Rebonato's [PLIGHT OF THE FORTUNE TELLERS: WHY WE NEED TO MANAGE FINANCIAL RISK DIFFERENTLY](#) (Princeton University Press, 2007). The writing is stylish, learned, and non-mathematical, but the book may be difficult going for those new to finance. For a short, breezy, but worthwhile account, see the article by Joe Nocera in the NEW YORK TIMES MAGAZINE, "[Risk Mismanagement](#)," 4 January 2009.



was relying on the historical average and the historical volatility of returns. (I was also assuming, contrary to the evidence, that extreme returns were no more likely than the bell curve formula implies.) The number I used for volatility represented risk, my uncertainty of what returns will be. But I don't know what the volatility and the average return will be in the future. In short, there are two levels of uncertainty about future returns: We take volatility to be the measure of our uncertainty, but we're uncertain about the degree of our uncertainty, because we rely on history to estimate it, and we're uncertain what exactly it is that we're uncertain about.

And it gets worse. Recall the bell-like plot of annual returns. Most of the returns are in the big, bulky middle. But financial ruin skulks in the left-hand tail of the plot, representing very large negative returns. Recall also that I said that the tail is fat, meaning that really bad returns have occurred more often than we'd expect if the curve were a true bell curve. But, however fat the tail, we have historically observed very few really bad annual or even monthly returns on the stock market. And that, in turn, means that, even if we could be sure that the future would be just like the past—and we're sure of the contrary—we have too few observations to have any justifiable confidence in inferences we'd draw from these historical data about the frequency and size of these confounding returns. Moreover, whatever the idealized shape that the statisticians may fit to the historical data, it is only an ideal approximation to the irregular reality of bad investment returns as they punished us in the past and as they may punish us in the future. None of us will live long enough to see the many financial disasters required to match the smooth mathematical curves postulated by the statisticians.

Can We Do Better?

Have the financial economists and risk managers been utter fools?

Of course not. I have not been suggesting that economists (at least the ones who specialize in finance) think *only* that risk equals the volatility of returns, and that they therefore don't understand risk. Neither proposition is true; well, maybe the latter, but only because the subject is deep, not because economists are simple-minded. The statistical approach to risk has indeed been fruitful, but mainly when applied to short spans of time. Measures of volatility tend to be stable over these intervals. We have lots of data for short intervals (seconds, minutes, hours, days, weeks, and to some extent, months). I have, myself, made ample use of the mechanical tools that the financial engineers have built using the statistical approach. But these tools do underestimate extreme risks, and they may not be applicable when there is an abrupt regime change in the larger economy.



How, then, should we estimate investment risk and manage our wealth around it? This is a larger topic to which I cannot do justice in this short essay. If risk were not a consideration, successful investing would be easy. Those who select individual stocks and bonds, whether private investors or professional advisors, can reasonably find comfort in Graham's margin of safety, but this isn't applicable to the full panoply of modern investment instruments, including mutual funds and options. It can also lead to over-engineered portfolios, like a nineteenth-century stone bridge that can safely carry a far heavier load of traffic than will ever actually pass over it. We cannot turn back the leaves of the calendar to the simpler financial world of 1950, nor should we. Alternatively, we can use the high-tech tools of the financial engineers, as long as we recognize that these are not well suited to long-term decisions, and that they may underestimate the likelihood of extreme events.

For longer-term planning, I prefer to be informed by the historical variability of returns, but I know that through them, I see the future only darkly. I continue to find useful the kind of mathematical analysis that I presented earlier, but I always remember that the results underestimate the chances of very bad investment results. I think of the estimated value corresponding to, say, a 1-in-20 chance of investments turning bad as an upper limit; the probability is likely greater than 1-in-20 that things will be that bad or worse.⁷ I also believe that it is necessary to take into consideration other kinds of risk besides volatility, pure and simple. Because the observed volatility of investment assets is a manifestation of underlying economic factors, we ought to consider the risks associated with those factors. And because risks are not constant at all times—that's why we have so much difficulty estimating future risks—we should take into account the prices and cash flows of our investments, as Graham did. It only stands to reason that if the prices of the investments are high when compared with expected cash flows, then the risk is greater than it would be otherwise. Mind you, this does *not* mean that high prices relative to cash flows portend a crash; rather, it means that the probability of a crash has increased. It also, for that matter, means that the probability of a modest decline in returns has increased. A forecast of increased risk is not the same as a forecast of lower returns.

Objectivity and Subjectivity

My purpose in this essay has been to begin to define and negotiate the concept of investment risk, not to give false assurances. I repeat, the bears haven't signed a treaty. All the same, despite the hard times that we are facing right now, one has to remember that such times are unusual, very unusual. Even though, in the words of the Primo Levi story, we now "taste bear meat," there are

⁷ Rebonato proposes a much more sophisticated way of addressing this issue.



reasons for hope.⁸ The U.S. government will not soon be taken over by a Communist cadre, nor will all our means of production be permanently nationalized by socialists. The U.S. has one of the younger populations among the developed countries, which means that it will continue to be highly productive. It is not time to hoard canned goods, buy guns, and move to a bunker in Idaho or Arkansas. If investment prices depend upon the outlook for the economy, we can at least say that the worst possible outcomes are not in sight.

So, having not given you false assurances, I've given you real assurances. That's because I don't know how, having considered what I just told you, you now *perceive* investment risk, and I don't want this essay to alarm you unduly. It's one thing to calculate or estimate the size of risk. It's another to judge whether it is terrifyingly large or comfortingly small, and that judgment depends on both your circumstances (that is, your wealth and your expected income and expenses), and your emotional composure. For example, if you are already enormously wealthy, with no great commitments for your money (such as to a charitable foundation), and of sanguine temperament, then you may easily weather a halving of your wealth.

The elegance of the definition of risk as volatility lies both in its making risk a matter of statistical calculation and in its separation of risk from the concept of *risk tolerance*. In modern financial theory, we first treat risk as an autochthonous idea, then we define risk tolerance as the psychology of risk. And to understand risk tolerance, we need first to explore the relationship between risk and return.

Conclusion

Investment risk should be viewed as a matter of probabilities, something that may be quantifiable, and that some amount of risk, even risk of extreme outcomes, is ever present. But I have also made three other points, namely, that it is useful to identify investment risk with the volatility of returns, second, that some common inferences from this identification are incorrect, and third, that a little statistical learning can lead to dangerous underestimation of investment risk.

Even when risk isn't quantifiable because of the vast unknowability of the future, we can nonetheless reason about it as if it were. Because risk is ever present for any investment, you should never trust anyone who implies that he or she can select investments that have no risk whatsoever. (Besides, the law forbids us advisers to make an explicit statement to that effect.) But not all investments have the same risks, and for any single investment, different outcomes have different likelihoods. For some investments, the chances of especially bad outcomes may be imponderably tiny.

⁸ Primo Levi, "[Bear Meat](#)," translated by Alessandra Bastagli, *THE NEW YORKER*, January 8, 2007.



I have played fast and loose with the kinds of investments I used to make my points: bonds, stocks, and markets. These all exhibit volatility of returns, but to different degrees, with different distribution curves and tails, and for different reasons. In particular, there is an important distinction to be made between the risk of individual investment securities and the risk of amalgams or aggregations of securities, which is what mutual funds and markets are. To understand this distinction, we need to explore the concept of diversification.

**Adam Jared Apt, CFA, is a financial advisor and the owner of Peabody River Asset Management, based in Cambridge, MA.*

www.advisorperspectives.com

For a free subscription to the Advisor Perspectives newsletter, visit:
<http://www.advisorperspectives.com/subscribers/subscribe.php>