Aaron Brown

Risk Manager Without Portfolio

You really should have some Kelly with your Markowitz

A couple of issues back, I griped about risk management textbooks leaving out a crucial empirical fact about Value at Risk, which is one of the quantitative foundations of the field. I'm going to bookend that column, not by praising textbooks, but by griping about something they put in them. I see too many textbooks that discuss portfolio management instead of risk management, and the confusion extends to many people – perhaps even the majority – in quantitative finance.

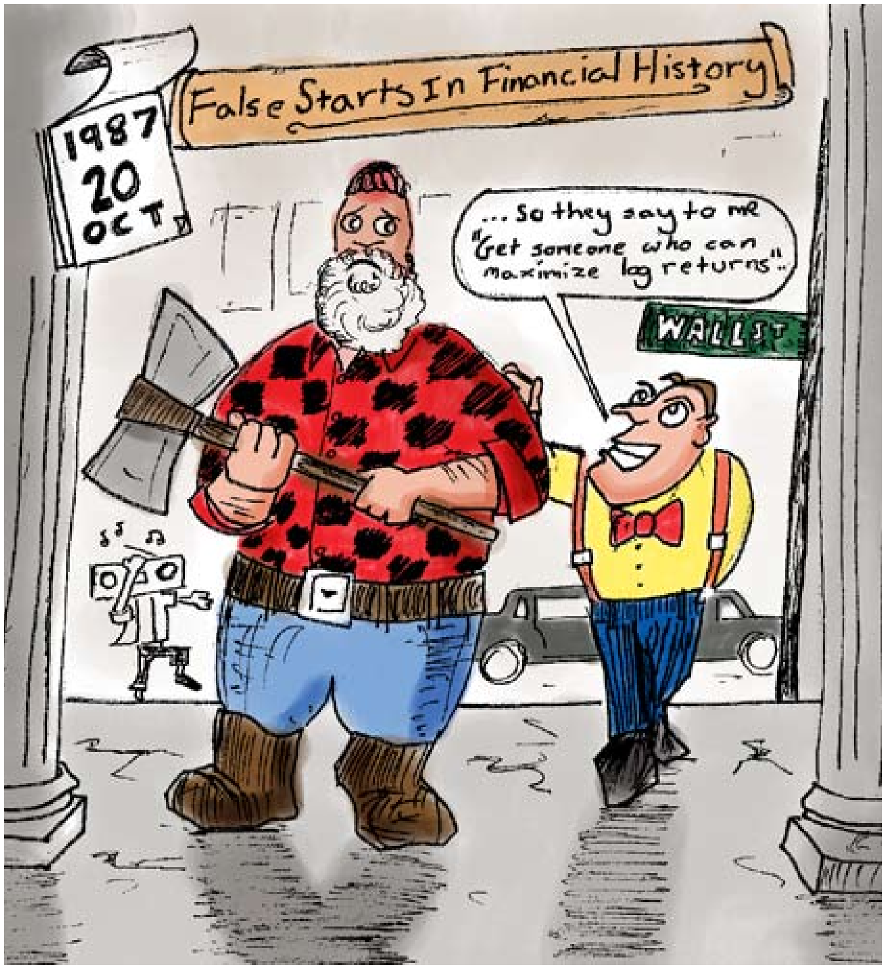

Let's start with some historical context. Harry Markowitz created the modern field of portfolio management with his 1952 *Journal of Finance* article, "Portfolio Selection." The problem he attacked is how to weight investments among a number of assets, such that the resulting portfolio probability distribution maximizes a utility function. This was a pathbreaking approach to an old problem, and it ushered in tremendous theoretical and empirical progress in finance, plus practical changes in the way people manage money. The revolution continues to this day.

Thirty-five years later, there were some quants working on Wall Street, trained in this theory, plus descendant ideas like the Capital Asset Pricing Model and Black–Scholes–Merton option pricing formula. The Crash of 1987 revealed an important problem that portfolio theory could not address. The field of risk management was invented in response. By 1993, there was consensus on the basic theory, and, like Markowitz's work, it stimulated tremendous progress in finance. Risk management is complementary to portfolio management; it does not replace it.

To see the difference, consider the following portfolio. Every quarter its mean return is an independent random draw from a Normal distribution with mean 2 percent and standard deviation 2 percent; and its volatility is an independent random draw from an exponential distribution with mean 4 percent. The return is then drawn from a Normal distribution with the specified mean and volatility. This results in a mean annual excess return of 8 percent (compounded quarterly) with a standard deviation of 12 percent.

Judging from these statistics, 8 percent alpha (since it is uncorrelated with the market, or anything else) with 12 percent standard deviation would make it a viable hedge fund. People get odd ideas about hedge fund performance due to reporting that seldom distinguishes among funds run at low volatility, such as 4 percent per year, that rarely have negative quarters and funds run at 25 percent or higher annual volatilities that try for much higher returns but lose money reasonably often. Another problem is overattention on a few superstar funds, almost all of which are closed and have tight capacity constraints, along with funds that have just finished hot three-year runs. A long-term track record of 8 percent alpha in a high-capacity, open fund with institutional transparency and controls, run at a 12 percent annual standard deviation, would make a very competitive product.

Over a 30-year horizon, this fund has less than 0.01 percent chance of returning less than the risk-free rate and a 99 percent chance of beating the risk-free rate by 3 percent per year or more. Its median performance returns 11.2 times what you would have from investing in the risk-free rate, and there's a significant upside – you have a 14 percent chance of doing more than twice that well. Of course, its zero correlation with the market means it is even more attractive combined with other investments. Overall, we have to judge this a big portfolio management success.

It is, however, a risk management failure. Why? It went out of business long before delivering those attractive, low-risk, long-term returns. It never got a chance to demonstrate the long-term track record that would attract investment. Studies of hedge fund survival identify drawdowns as the best predictor of funds that liquidate. While the data have huge problems of self-selection and inadequate and inconsistent reference data, it seems that a drawdown equal to one annual standard deviation will kill about 50 percent of funds, and two annual standard deviations is seldom survivable, except by the best-established funds (the past 18 months may rewrite some of these statistics, as investors have learned, painfully, to lower their expectations).

The *expected* maximum drawdown for this fund is more than two annual standard deviations. It spends 14 percent of its time in drawdown greater than one annual standard deviation, and 4 percent in drawdown greater than two annual standard deviations. Even before you get to those levels, you can trigger a cascade of problems. A period of poor returns causes some redemptions. Prime brokers and counterparties look at declines in assets under management (AUM), which combines the effect of drawdowns and redemptions. They may tighten margin terms which work along with the decline in AUM to force position reductions, and those required reductions are likely to be known to others. Worse, funds with similar strategies may be in similar positions, making the same reductions. That means these position reductions may be made on unfavorable terms, which further hurts returns, bringing us back to where we started, poised for another turn of the spiral. Even if the strategy rallies, the reduced position sizes mean the fund will not make money back as fast as it lost it. And all these people, investors, prime brokers, counterparties, and opportunistic traders in the market are trying to guess if any of the others will pull the plug on the fund; in which case, they want to be sure to pull it first. This can make for a nervous business.

Things generally work best with portfolio managers concentrating on putting on positions to give the best distribution of returns, and a separate risk manager concentrating on the long-term cumulative effect

Note that this takes place even in our gentle toy model, without such real-world risks as autocorrelation, dependence between mean return and volatility, market-wide disruptions, operational errors, and unknown and changing underlying parameters; and with only very mild fat tails. Risk management is the quantitative field that considers not the one-period expected utility a fund delivers, but its multiperiod performance considering the interaction between fund performance and the rest of the world. It's not just about minimizing drawdown – that's only a small part of risk management. The tagline version is that the risk manager makes sure the fund survives long enough to realize the portfolio manager's Sharpe ratio.

Be careful to distinguish between people and jobs. Portfolio managers knew all about the dangers of drawdown before risk management was invented, just as people who managed money knew about diversification before Markowitz. Good portfolio managers today automatically consider risk management issues when making investment decisions. It is entirely possible to run a good modern portfolio without giving anyone the title of risk manager. But just as Markowitz gave a theoretical underpinning and a set of mathematical tools to exploit diversification, the field of risk management is a theory and set of tools for managing the long-term outcome of a series of short-term decisions. You don't need the formal field to make good risk decisions; some people do a good job intuitively, or with traditional rules of thumb. Conversely, you can do a terrible job of risk management using all the shiny, fancy theory and tools of the modern field. Still, I believe that things generally work best with portfolio managers concentrating on putting on positions to give the best distribution of returns, and a separate risk manager concentrating on the long-term cumulative effect of the series of decisions.

I'm not just talking about portfolio management, by the way. To a risk manager, any risk-taking project looks like a portfolio. A pharmaceutical company researching a new drug, a movie studio filming a romantic comedy, an appliance company managing a line of refrigerators – they're all probability distributions that interact with the world. The risk manager doesn't tell a portfolio manager which positions to take, nor does he/she suggest chemical compounds, rewrite the movie ending, or choose colors for refrigerators. His/her job is to identify the important interactions with the world, model the outcome distributions of the projects jointly with those interactions, and estimate the long-term results of alternative strategies. His/her main output is optimal risk level: what volatility the portfolio should be run at, and whether to raise or lower the risk of business projects by things like changing the budget or schedule, taking on more or fewer projects, looking for risk-sharing partners, or taking more or less predictable approaches.

Let's go back to Markowitz. A simple version of his problem assumes a known set of assets, with a multivariate Normal excess return dis-

>

tribution over a fixed horizon that has a known mean vector and covariance matrix. If we want to maximize the Sharpe ratio, we invert the covariance matrix and multiply it by the mean vector to get a vector of optimal portfolio weights.

Of course, real portfolio problems are tougher; we generally have to choose the universe of assets ourselves, we have to estimate distributions and parameters, and we can have more complicated objectives. Plus, we have to consider things like transaction costs, leverage terms and dangers, liquidity, short availability, taxes, investment restrictions, and so on. But there are three aspects of the toy problem that are close to real portfolio management. First is that we consider a specific horizon, both for distribution estimation and objective function definition. The assumption is that we'll redo the analysis at the horizon and choose a new portfolio for the next period. This assumption is subject to some constraints; some assets are expensive or impossible to trade, and some objectives are defined over multiple periods, so in some cases we expand to consider more periods, but it's still a finite-period problem. Second is that our answer is only a relative weighting among the assets; there is nothing to tell us the absolute size of investment. Again, this may be subject to some size constraints on specific assets or the portfolio, or limitations like no short-selling or a leverage ceiling, but it's still fair to say that portfolio management concentrates on relative allocations. Third, we assume some kind of objective function, which could be a utility function.

The simple version of the risk management problem is that you're going to make an infinite series of bets. Each bet will have a known probability distribution at the time you make it, but you don't know today what opportunities will be offered in the future. You can make each bet in any size you want, up to your total wealth at the time.

The Markowitz of risk management was John Kelly, who proved in 1956 that the answer to the problem above is to choose the bet size that maximizes the expected value of the logarithm of wealth. Note that there's no time horizon, no relative weightings, and no utility function, all of which distinguish it from the portfolio management problem. Instead of relative weights, it gives us the absolute amount of risk to take.

A simple way to remember the difference is that portfolio managers are concerned with a linear approximation with mean excess return divided by standard deviation (Sharpe ratio), while risk managers' linear approximation is mean excess return divided by variance (Kelly criterion). These are in different units, different dimensions. The Sharpe ratio is invariant in bet size, but changes with time horizon. The Kelly criterion changes with bet size, but is invariant in time horizon.

Another quantitative distinction is that the portfolio manager's first concern is expected value, while the risk manager starts from the worst-case outcome. A portfolio manager who makes negative expected return bets will always fail in the long run; a portfolio manager who makes positive expected return bets, even if he screws up the diversification, at least might succeed. Even if he gets fired as portfolio manager, he will be a valuable submanager or analyst. Risk management formulas, including the Kelly criterion, turn out to be highly sensitive to the worst possible outcome, even if it has very small probability. If you stay in business long enough, sooner or later you'll experience the worst case of some bet, so you have to make sure that it's a survivable event. Only after that can you think about expected profit, because expected profit is only useful to someone around to enjoy it.

Real risk management problems are more complex. Horizons aren't infinite, and bet outcome distributions are not known with certainty. Bets can be overlapping in time and correlated, and may not be available in all desired sizes. Wealth will generally be neither the goal nor the constraint. Goals don't present too much problem, even when you have to align interests of different parties, but real-world constraints turn out to be strongly path dependent, highly multidimensional, uncertain, and nonconstant.

The distinction between portfolio and risk management is sharp only in theory. Portfolio managers optimize simultaneous bets and risk managers optimize nonoverlapping independent ones, but both have things to say about overlapping and dependent bets. Portfolio managers have responsibility for single horizons and risk managers consider infinite repetition; finite multiple period problems draw on insights from both disciplines. Asset management companies often have many portfolio managers whose portfolios are combined in different ways in different products; this can be regarded as either a portfolio management problem with portfolios instead of assets as the underlying units, or a risk management problem in optimal sizing and combination of portfolios. In practice, it requires the tools of both disciplines.

However, just as good fences make good neighbors, it's important to avoid conflating the two optimizations. Doing both at once requires thinking about not only the probability distribution of outcome of any combination of outcomes you might take, but also on the distribution of all future distributions. I doubt such a thing is possible, even in theory; I think it embeds logical contradictions, and I know it is not a useful paradigm in practice.

A practical reason for separation of the problems is that they are quantitatively different. Portfolio management problems tend to be high-dimensional and involve gathering and analyzing large amounts of data. Risk management problems tend to be low dimensional, with much smaller data requirements. Portfolio management is almost forced to adopt some parametric techniques, although these are usually modified for robustness. Risk management is primarily nonparametric.

Harry Markowitz was both the academic who wrote the seminal paper in his field, and a practitioner for 58 years and counting. John Kelly wrote the risk management paper, and did a lot of other good stuff, but never exploited the criterion, and unfortunately died of a stroke at age 41. It was Ed Thorp who pioneered the practice of risk management, building on Kelly's work, and incidentally proved that a portfolio manager can be a risk manager at the same time. The modern independent risk manager, however, was invented between 1987 and 1993, by people trained in portfolio management, aware of Kelly and Thorp, but also bringing their own perspective to the problem.

That story is contained in the 10 steps of quant finance enlightenment:

1. I know nothing about markets.

2. I learned basic efficient market theory, which strips away tremendous nonsense and leaves me with a beautiful theoretical picture of the world that fits remarkably well. Still, there are anomalies. But from the classroom, there's no way to tell the difference between an exploitable anomaly and a mirage caused by coincidence, data error, or imperfect theory.

3. I've moved to Wall Street and it's now easy to identify the exploitable anomalies. I can see why prices are the way they are, and how a simple application of mathematics with financial theory can give me free money. There are lots of these ideas around. There must be some catch. I'd better quadruple check my results and start very cautiously. It can't be this easy.

4. It is this easy. I'm going to be infinitely rich.

5. Okay, it's not that easy. Yes, exploitable opportunities are easy to find. But they come in two types. Some have a very high probability of winning, but require a lot of time and effort for each bet, and have limited size. The Sharpe ratio is off the charts, but you have to manage carefully to get a good return on your time and to keep your assets deployed. They other type is consistent and available in size, but has only moderate Sharpe ratio. You need to manage patiently, without errors, in order to make a very large profit safely over a moderately long period of time. It appears that the bet sizing problem is as important, or maybe more important, than the choice of which opportunities to exploit. Let me review my Thorp and Kelly.

6. Okay, I've got it nailed – this is easy. I'm going to be very rich.

7. How come no one told me that things like the Crash of 1987 happen?

8. Financial modeling is impossible if we start with the assumption that there is a fully specified probability distribution over future events, even if you let that distribution be unknown. We need to allow for the possibility that prices are undefined sometimes. The advantage is that our portfolio management just got a whole lot easier; we can find consistent multivariate distributions to feed to our portfolio optimization algorithms that are simpler and work much better than the old ones that tried to account for every possible outcome. We can reconcile market-implied and actual probabilities. We can build models that give both realistic dynamics for price evolution and accurate cross-sectional prices for securities and derivatives. One disadvantage is that we expect our model to fail with a specified small frequency. Another disadvantage is that the distributions that work best for portfolio optimization are different from the ones that work best for risk management.

9. Wow! When you figure this stuff out carefully, it turns out to be much more important to take maximum advantage of your best opportunities than to avoid drawdowns and other bad events. No amount of conservatism can reduce the downside below a certain level, but you can increase the upside dramatically, which can give you the cushion to survive the inevitable knocks. Defensive risk management is still very important, but, as George Washington put it, "Offensive operations, often times, is the surest, if not the only means of defense." Another insight is the importance of paths.

10. I think I've got this licked. I've found some great opportunities and I know the math to optimize both relative weights within each period and bet sizes over periods. I've got a reliable Value at Risk with a successful out-of-sample backtest, solid stress tests, and contingency plans. I've modeled my portfolio problem, and also my multiperiod objective function and constraints, and run simulations to determine the best strategy. Now, if there are just no more surprises...

Once upon a time, there were three little piggies whose parents gave them $1 each to make their fortunes in the market. Piggy number one spent all his time studying portfolio management, and got so good that he could win even money bets 90 percent of the time. He was too busy to learn risk management, however, so he bet all his money each time. After 10 bets, he had $1,024, but he lost everything on the 11th bet and was turned into bacon.

The second little piggy devoted half his time to risk management. He realized that the first piggy had actually been lucky; only about one-third of the time will someone following that strategy survive for 10 flips. And sooner or later, the first piggy had to fail. Since the second little piggy didn't have as much time for portfolio management, he learned to win even money bets only 60 percent of the time. But he knew enough risk management to compute that he had only a 1 percent chance of losing five such bets in a row. So, he bet $0.20 each time.

The third little piggy spent all his time on risk management. He only learned how to win even money bets 52 percent of the time. But he knew his Kelly, and so he bet 4 percent of his wealth on each flip.

After 10 bets, he had $1,024, but he lost everything on the 11th bet and was turned into bacon

The second little piggy made an average of $0.04 on each bet, and had average luck to run his bankroll up to $5 after 100 bets. He sneered at the third little piggy, who also had average luck but whose total wealth was only $1.08. After 1,000 bets, the second little piggy was up to $41, while the third little piggy had only $2.23. But something strange happened on the way to 10,000 bets. The second little piggy had only $401, while the third little piggy had $2,987! At 30,000 bets, the second little piggy was up to $1,201, passing the first little piggy's high water-mark, but the third little piggy was the richest pig in the world, at $27 billion. It never got to 40,000 bets, because before that happened the third little piggy had all the money in the world, so there was no one left to bet with.

The moral of the story is, make sure your risk manager is the right pig. W