Aaron Brown

The Theory of Risk Management

A look at the core ideas of what we do through a brief history of finance as an academic subject

I am often asked to recommend a book that explains financial risk management. While there are several excellent handbooks that explain how it is done, oddly enough there is no good book or article that covers the basic theory. This was not an important problem when the field was small and everyone knew each other. But as the field ages and grows it's easy for new entrants, who were not involved in the initial development, to lose sight of the core ideas.

People have been worrying about risk for as long as there has been worry or risk. Financial institutions have long had legal staffs, regulatory compliance officers, credit officers and auditors to reduce risk. General economic activity made use of a wide range of specialists to reduce physical risk, from civil engineers and safety inspectors to security guards. Insurance actuaries have been measuring risk for centuries.

None of this has any relation to modern financial risk management. The fields are entirely distinct. They have different goals, assumptions and methodologies. The basic words have different meanings. Traditionally, 'risk' meant the chance of something bad happening and 'risk management' meant reducing that chance. In modern financial risk management, 'risk' is any deviation from expectation and 'management' refers to managing the business using risk as a tool. It's not the risk that's managed, it's the business.

The fundamental principle of risk management revealed itself like a thunderclap 15 years ago, apparently to many people at about the same time. It gave rise to a burst of creativity that revolutionized finance and indirectly the whole economy. Its implications are still being worked out. To understand what happened, we need a brief review of the history of finance as a field of study.

A brief history of finance

Until the 1950s, academic finance was a purely descriptive field, like biology before Darwin. Students learned what a trade letter of credit was and what documentation was needed for a bond underwriting. It was more a branch of law than business, students went into finance if they didn't have the quant skills for accounting. Then came an extraordinary generation of professors: Franco Modigliani and Merton Miller in corporate finance; Harry Markowitz, William Sharpe and Eugene Fama in portfolio management; Fischer Black, Myron Scholes and Robert Merton in derivatives; plus others of equal stature. In less than 20 years we went from pure description to having consensus models for the most important areas of finance: capital structure, capital asset pricing and price dynamics.

Unfortunately, the models were not very satisfactory. They were immensely important, they cleared away centuries of nonsense and put the field on a firm theoretical foundation. But the Modigliani-Miller Theorem was not supported by convincing evidence and, in any event, it said

nothing mattered. The Black-Scholes model could tell us an option price, but only if we knew the future volatility of the underlying. This seemed to be a step backward. We didn't know either one, but at least the price was observable and in the present, while the volatility was unobservable and in the future. The capital asset pricing model was even worse. It could tell us one observable present variable, but only if we fed in three future variables, two of which were unobservable. An enormous effort went into testing the model empirically, with inconclusive results. None of the models were any good for making money, or making any financial decision, because the input data were too hard to estimate, and they said nothing mattered anyway.

Academics on the Street

Nevertheless, the first generation of students trained in these models went to Wall Street and got very rich. In 1984 a single desk, the Salomon Brothers mortgage trading desk, made more money by far than any entire financial institution in history. And that was just the beginning. The new finance, despite apparent uselessness, fueled undreamed of profits. The models were all zero-sum, but the application of the models turned out to have spectacular positive sum. Not just Wall Street, but the entire economy generated unprecedented profits as a result of scientific finance.

The advances fell into three main areas. Leveraged buyouts and related transactions proved that capital structure definitely did matter. By shedding assets and taking on debt, businesses reversed 15 years of negative real returns on equities and ushered in the greatest and longest bull market in stock market history. Derivative trading quickly eclipsed trading in underlyings, resulting in enormous trading profits and spectacular feats of financial engineering. Statistics-based models created the mortgage-security market and deciphered basic fixed-income dynamics.

Innovation, as always, led to disasters. By the late 1980s corporate raiders, LBOs and junk bonds were being blamed for all kinds of social ills, and some of the brightest lights in the field were facing prison. Derivatives were blamed for

In 1984 a single desk, the Salomon Brothers mortgage trading desk, made more money by far than any entire financial institution in history. And that was just the beginning

the largest one day crash in the history of the stock market and farther from the headlines, Andy Krieger's FX option marking problems highlighted the problems of managing the high-powered new techniques. Mortgage securities played a large role in the horrendously expensive S&L debacle and statistical portfolio techniques were blamed for the overexposure of major banks to credit problems in Latin America.

The fundamental principle of risk management

Disaster, as always, led to short-lived agitation to restrict innovation. But at two institutions top management decided that running stable businesses in the new financial order required more, not less, innovation. Bankers Trust encouraged quants to investigate return on risk-adjusted capital, which led to the modern concept of Economic Capital, and JP Morgan came up with Value-at-Risk. Both of these were designed as tools for the CEO to manage the businesses, not as tools to change the level of risk.

Also in the late 1980s a number of quants had suddenly noticed that an important change had taken place:
The risk of a business can be measured far more precisely than its expected return.

This is exactly the opposite of conventional wisdom in statistics. For a random variable X, the risk is usually measured by the expected value of X^2 or X to some higher power. These higher moments are generally harder to estimate than the expected value of X. For example, consider a distribution with mean μ and variance σ^2, except that there is some small probability p of an extreme value Y. It is rarely possible to rule out such small probability events either by observation or economic theory.

The new mean of the distribution, μ', is equal to $\mu + p(Y-\mu)$. The new variance of the distribution is $\sigma^2 + \Delta_\mu^2(1/p-1) - p\sigma^2$, where Δ_μ is defined as the change in μ, $p(Y-\mu)$. For small p, the change in σ^2 will be many times the change in μ. For higher order moments the formula is more complicated as it depends on all lower moments, but the formula for the change in the expected value of X^n contains the term $\Delta_\mu^n(1/p-1)^{n-1}$ which makes them even more unstable than the variance.

The annual real log return on a diversified US stock portfolio has a historical mean of about 4 per cent with a standard deviation of about 20 per cent. Suppose I hypothesize that there is a 0.0001 chance of a −99.96 log real return each year. There is no good theoretical reason for this not to be true, and the fact is has not been observed is small evidence against it, since we have only about 80 years of good stock market history and at most 350 years of any history. This would reduce the expected annual real return on the stock market from 4 per cent to 3 per cent, but the standard deviation would soar from 20 per cent to 102 per cent. So neither theory nor empirical evidence gives me any solid basis for estimating the standard deviation of stock returns. If the original skewness was zero with kurtosis of 3, the new distribution is −94 skewness with over 9,000 kurtosis.

But derivatives quants saw the world differently. First they noted that for the parameters above, without the extreme event, with 25 years of data the standard error of estimate for the mean is 4 per cent, so there is 100 per cent measurement error. The standard error of estimate for the standard deviation over that period is under 6 per cent, less than 30 per cent measurement error. But more important, if we take more

frequent measurements, it does not help the mean estimate at all while the standard error of the standard deviation depends only on the number of measurements, not the intervals between them. 260 daily observations tell us the standard deviation with a standard error smaller than 3,000 years of data for the mean, however often we measure. We can improve the standard deviation estimate even more by using high and low values, this information is useless for estimating the mean. Most important of all, a single option price tells us the market expectation of standard deviation with negligible error.

How the world changed

The 1980s saw dramatic increases in market liquidity, both in underlyings and derivatives. Data became easily available and techniques were invented to measure not just standard deviation, but subtle patterns of volatility by underlying price and tenor. During the same period, people continued to debate whether the expected real returns on stocks and bonds, the simplest financial instruments, were even positive. Quants got used to a world in which nothing useful was known about expected return, while volatility of return was precisely measurable.

During the same period, people continued to debate whether the expected real returns on stocks and bonds, the simplest financial instruments, were even positive

How does this get around the problem of a hypothetical low-probability extreme event? If the probability is constant and common to all financial instruments, or if its existence is unsuspected by any market participant, then volatility models ignore it. That's the businesslike thing to do, because in this situation you can't know anything useful about the risk, and you can't do anything useful to mitigate it. So risk is redefined to include only things we can know about and control. This is an important point. Traditional risk minimization includes all risks, which means financial risk management tools are only partially applicable. It also makes traditional risk minimization inherently unscientific. Its dictates cannot be objectively verified or studied by controlled experimentation.

Statistically inclined finance quants with mortgage backgrounds came to the same conclusion in a different way. It is often possible, especially in fixed-income, to know the possible outcomes of an investment. In this case the expected return is generally much more sensitive to the probabilities of the outcomes than the standard deviation is. Consider an even-money bet, for example. There are two outcomes, win $1 or lose $1. If p is the probability of winning the expected value is $2p-1$ dollars and the standard deviation is the square root of $p(1-p)$. At $p = 1/2$_ the elasticity of expected value with respect to p is infinite, while the elasticity of standard deviation is zero. When the expected return gets more than 2 standard deviations away from zero the elasticity of the standard deviation becomes larger than the elasticity of the expected return, but such situations are seldom encountered in practice. A bet at that level of statistical confidence is nearly an arbitrage.

More generally, it's possible to construct exact confidence intervals using resampling methods. These methods tell us nothing about the expected value. They avoid the problem of the low-probability extreme events by limiting prediction to within the range of historical values. We can get a precise estimate of the probability of, say, an annual return on the S&P 500 worse than the worst year in history; but we can't say anything about the conditional distribution of such a return. That's fatal for estimating the expected return of the S&P500, but poses no difficulty for risk measurements that do not go too far into the tail.

The business of risk management

Just because quantitative techniques are not very useful for measuring expected returns of economic decisions, doesn't mean we shouldn't be concerned with expected return. It just means we have to rely on the subjective judgment, generally non-quantitative, of expert business managers.

The first use people think of for precise risk measurement is preventing catastrophic losses. This idea underlies a lot of the funding for Risk Management departments. But modern financial risk management ignores a lot of potential catastrophes and isn't designed to prevent anything. We still rely on prescientific techniques to guard against total disaster: legal analysis, audits, regulatory compliance, stress testing, credit monitoring and knocking on wood. Risk management should be integrated with these fields, for the same reason a doctor should be cheerful and wish you good luck while prescribing a pill whose safety and efficacy were validated by large, double-blind clinical trials. Neither scientific risk management nor scientific medicine has progressed enough to despise traditional remedies.

It is often remarked that the institutions best-known for their risk management innovations have not avoided disaster. Bankers Trust and JP Morgan, the originators, no longer exist as independent institutions in part due to problems that better risk management should have avoided. Askins Capital Management, Long-Term Capital Management and Enron, the most aggressive risk management innovators, all crashed. Meanwhile my current employer Citigroup, is the largest and most profitable financial institution in the history of the world. It resulted from a combination of Salomon Brothers, famous for taking more uncalculated risk than anyone, and Citibank, famous for rejecting modern risk management (of course today the firm is among the forefront of risk management practice). This should convince anyone that risk management is not about

avoiding disaster.

The next thing people think of risk management is that it's useful for optimizing risk. Even experienced business managers have trouble making good decisions with respect to low-probability events. They don't happen enough to get much experience at them. Often low-probability events are either ignored or avoided at too high a cost. The cumulative effect of such misjudgments can be significant. It's also hard to aggregate risks qualitatively among different types and levels of business. Precise risk measurements can help business managers make better decisions.

But this is only half the story. Risk management is also an important strategic control system. At Bankers Trust and JP Morgan, we went back and forth between the businesses and top management, trying to come up with measures that made sense to both groups. The business implementations were necessarily highly specific, the measures that made sense to mortgage securities were different from the measures that made sense to equity underwriting or FX trading. Fortunately, we had quants who had come up through all of these businesses.

Top management was not primarily concerned with control or optimization of risk. They wanted to control the businesses. They wanted to allocate capital to the successful ones and starve the rest. They wanted to bring businesses along carefully, making sure each step was in the right direction, and that no rogue department would embarrass the firm. They needed general measures that applied to all businesses. More important, the measures could not be just measures, they had to be integrated with control parameters.

The CEO problem

Suppose you were the CEO of a modern multinational financial institution, without detailed knowledge of all the businesses of your firm. If your only concern is preventing some trader of a product you never heard of in a country you can't spell from bankrupting the firm, traditional limits work fine. But if the problem is attracting and motivating top individuals, nurturing risk-takers with capital and weeding out incompetents and reckless gamblers, quickly attacking promising new opportunities while nimbly sidestepping new risks, limits are no help at all. You need to track P&L closely, but that is not enough. It's easy to generate profits by taking low-probability risk, say by writing lots of out-of-the-money options. You can't estimate the risk of such a business by statistics on the P&L.

Risk management is about improving the business, not reducing the risk. In fact, a well-managed institution will generally take much more controlled risk than any badly-managed institution could afford

Modern risk management reports give you a metric for comparing businesses. You can tell which ones are demonstrably profitable, which ones are promising, which ones are doubtful and which ones are clearly bad. You can gradually reallocate capital for maximum value, keeping close track that businesses do not change in character. Not only do you have a better portfolio of businesses, but the businesses are individually better for constant precise feedback. Good managers shine when they have to explain their decisions, lucky incompetents are exposed.

Risk management is about improving the business, not reducing the risk. In fact, a well-managed institution will generally take much more controlled risk than any badly-managed institution could afford. The key is that the risk is productive and managed by the most talented individuals.

A side benefit of this is you have a much better idea of likely future events. This allows you to optimize your capital structure. Well-managed institutions get more liquidity while holding fewer low-return assets, more credit at lower cost and enjoy better price/earnings multiples in the stock market because management can demonstrate it's in control.

The key to all of this is risk measures that are accepted by the businesses and useful to top management. Developing them took lots of careful back and forth work (something that is not being done for Operational Risk, which troubles me about the field). Economic Capital and VaR work in both domains. That point is missed by many people without practical risk management experience who champion alternative measures that were tried and rejected ten years ago. Of course there is much useful work to be done, and it's entirely possible new measures will supplant the old. But it will be done by people who understand the basic theory of risk management.