Summary

The history of quantitative risk management is clear when viewed through the lens of Modern Portfolio Theory. Many of the most important developments in risk management and many of our most challenging problems are readily apparent within this framework.

About the Author

Michael B. Miller studied economics at the American University of Paris and the University of Oxford before starting a career in risk management. He is a former member of GARP's FRM Exam Committee, the author of *Mathematics and Statistics for Financial Risk Management*, an adjunct professor at Rutgers Business School and CEO of Northstar Risk Corp.

The History of Quantitative Risk Management and Modern Portfolio Theory[1]

In 1952, *The Journal of Finance* published "Portfolio Selection" by Harry Markowitz. The article introduced the world to Modern Portfolio Theory (MPT). For this and related work, Markowitz would go on to win the Nobel Prize in Economics. The central insight of MPT is simple, elegant, and difficult to dispute. Putting this simple idea into practice, however, requires us to make three basic assumptions. The history of quantitative risk management can be viewed as a continuing effort to refine these assumptions.

The central insight of MPT is that investors are trying to get the highest returns with the least amount of risk. Given two portfolios with the same level of risk, but different expected returns, a rational investor will prefer the portfolio with the higher expected return. Similarly, given two portfolios with the same expected return, but different risk levels, a rational investor will prefer the less risky

[1] This paper is part of a series of papers by Michael B. Miller, which can be found at www.risk256.com. Please do not reproduce without permission. To request permission, or to provide feedback, you can contact the author, mike@risk256.com. This version of the paper was last updated April 29, 2014.

portfolio. That this seems obvious — that it seems natural to frame investing in terms of risk and return — is a testament to the profound impact that MPT has had on finance and risk management.

As stated, the central insight of MPT is hard to dispute. All else being equal, rational investors will try to maximize returns and minimize risk. The real problems arise when we try to put this simple statement into practice. Markowitz himself made two key assumptions:

1. That risk can be equated to standard deviation.
2. That the relationship between securities in a portfolio can be fully described by correlation.

To put these first two assumptions into practice, we typically make a third assumption:

3. That returns, standard deviation, and correlation of assets can be estimated based on past returns.

There are many other technical assumptions that we need to make such as the ability to borrow at the risk free rate, and that there are no transaction costs, but these are the three big ones. These are the three assumptions that I believe have defined quantitative risk management ever since. We review each in turn.

Standard Deviation = Risk

By 1952, standard deviation had already had a long history in statistics. It is an extremely popular measure of volatility in a number of fields and very easy to work with. It is not surprising that Markowitz chose the standard deviation of returns as a proxy for portfolio risk. As a first order approximation, standard deviation is actually not a bad choice, and standard deviation is still widely used when characterizing the risk of securities and portfolios.

One problem with equating risk and standard deviation in finance is that the standard deviation of returns is rarely stable over time. If we were to attempt to characterize the risk of a portfolio using standard deviation, we would have to first ask ourselves which standard deviation we wanted to measure. Is it the expected standard deviation tomorrow? The average standard deviation next month? Next year? Even though the standard deviation of securities tends to vary over time, the variation is far from random. For many securities, the standard deviation of returns exhibits strong serial correlation and mean reversion. If the standard deviation of a security is high today, it is likely to remain high in the near term, as it slowly drifts back towards its long-run average. In the 1980s, awareness of these types of patterns led economists, starting with Robert Engle in 1982, to propose a number of models of time-varying standard deviation. Since then, generalized autoregressive conditional heteroscedasticity (GARCH) models have become extremely popular in finance. Just as the original MPT model imagined investors attempting to minimize the risk of their portfolio for a given expected return, it is logical to expect investors will seek to minimize the volatility of the standard deviation of their portfolios. All else being equal, investors will likely prefer a portfolio with a more stable and predictable standard deviation.

Not only does the standard deviation of returns tend to vary over time, but occasionally we observe extremely large returns or jumps. Robert Merton first introduced the mathematics of jump-diffusion processes to finance in 1976 in the context of options pricing models. The applicability to risk management was obvious, and jump-diffusion models are now widely used in financial risk management. In many markets, jumps are asymmetric. For example, in equity markets large negative returns are much more likely than large positive returns. There are no "up crashes". If a portfolio contains assets that exhibit negative jumps, failing to model those jumps could lead to a severe underestimation of risk. All else being equal, investors will prefer a portfolio with fewer negative jumps.

Even in the absence of jumps and time varying volatility, standard deviation only begins to describe how returns vary about their mean. Standard deviation is the square-root of variance, which is the second central moment of a distribution. Higher order moments get little attention in many other fields, but are extremely important in risk management. Today, skewness and kurtosis — the third and fourth standardized central moments, respectively — are standard metrics in the risk manager's tool kit. All else being equal, investors tend to consider portfolios with negative skew and higher kurtosis to be more risky.

In quantitative risk management we tend to focus on the distribution of returns, but portfolios can vary in other important ways as well. It can be difficult to fully characterize the liquidity of a portfolio, but liquidity is certainly an important risk characteristic. All else being equal, investors will prefer a more liquid portfolio.

Beyond these refinements there are a whole host of technical and psychological issues that may be important for actual investor, including tax efficiency, transparency, and the frequency of losses.

Time varying risk, jumps, higher order moments, and liquidity continue to be a major focus of quantitative risk management. The theoretical foundation is very solid, but integrating these refinements into a daily risk management process can still pose a challenge.

Portfolio Risk = Correlation

Just as it is easy to understand why Markowitz used standard deviation to characterize the risk of a portfolio, it is easy to understand why he chose correlation to describe the relationship between securities. If we know the standard deviation of individual securities and their correlation to each other then we can easily determine the standard deviation of any portfolio containing those securities. If risk is standard deviation, then correlation is all that you need to describe the relationship between securities.

The limitations of correlation are by now well known, but how to address those limitations is far from being a settled matter. Correlation and covariance measure the average degree of association of a pair of random variables. If two securities are highly correlated then they tend to have above average and below average returns at the same time. But what if below average returns are more likely to occur together than above average returns? What we often see in finance is that a pair of variables will have a higher association in one direction than in another. That equities are more likely to have large negative returns at the same time is one reason why the distribution of stock market returns tends to be negatively skewed. When the association of two random variables is symmetric, we say that the joint distribution of the two variables is elliptical. When the association between two random variables is not symmetric — as is typical of equity returns — we say that the joint distribution is non-elliptical. The problem with correlation is that it cannot differentiate between elliptical and non-elliptical distributions (or between different types of non-elliptical distributions). What I outline next are three possible parametric approaches to moving beyond correlation: copulas; higher order cross moments; factor analysis and stress testing. For the purposes of risk management, at some level all three approaches are equivalent; all three methods can describe the type of extreme portfolio returns that are the result of non-elliptical joint distributions. From an implementation standpoint and from an intelligibility standpoint, these three approaches are very different. There are pros and cons to each.

A copula combines the cumulative distribution of two or more random variables to produce a single joint distribution. Copulas are extremely flexible and can produce a wide range of non-elliptical distributions. The mathematics behind copulas can be extremely elegant, but also extremely complex. Because the math is so complex, copulas are very rarely covered in entry-level or even intermediate-level statistics textbooks. While risk managers are increasingly familiar with the concept of copulas, few fully understand the mechanics. Also, the complexity and data requirements associated with copulas quickly grow as we add more variables. In practice copulas are likely to be limited to describing the relationship of a small number of random variables.

Higher order cross moments are a logical extension of covariance. Covariance is the second cross central moment, and the covariance of a random variable with itself is simply the variance of that variable. Similarly, coskewness is the third standardized cross central moment, and the coskewness of a random variable with itself is simply the skewness of that variable. Kurtosis is likewise the fourth standardized cross central moment. Just as a single variable that exhibits negative skewness is more likely to generate extreme negative deviations from its mean, two variable that exhibit negative coskewness are more likely to generate extreme negative deviations at the same time, than they are to exhibit extreme positive deviations at the same time. Calculating coskewness and cokurtosis is no more difficult than calculating covariance. All three can be easily calculated in a spreadsheet (for a review of the mathematics of moments and cross moments, see Miller, 2012). One potential drawback of higher order cross moments is that the number of non-trivial cross moments is large for even a small number of variables. Two variables have one non-trivial covariance (e.g. covariance of X and Y), but two non-trivial coskewness statistics (e.g. XXY and XYY), and three non-trivial cokurtosis statistics (e.g. XXXY, XXYY, and XYYY). For ten random variables there are 45 non-trivial covariances, 210 non-trivial coskewnesses, and 705 non-trivial cokurtosis statistics. In finance we are unlikely to have enough data to fully model all the higher order cross moments in a large portfolio. As with copulas, higher order cross moments are most likely to be limited to models of a few variables.

Though not often thought of this way, factor models can be a very powerful and straightforward way of describing complex relationships between variables. As used here a factor model refers to a model where a set of random variables are a function of more fundamental random variables or factors. A very simple example is an equity portfolio where each equity in a portfolio is a linear combination of a shared market variable and an idiosyncratic random variable. By changing the weight on the market and idiosyncratic components of the equity returns we can change the correlation between the equities in the portfolio (all of the weight on the market component corresponds to perfect correlation; all of the weight on the idiosyncratic component corresponds to no correlation). By changing how we model the factors we can describe a wide range of relationships between two variables. For example, pretend we have two equities in a portfolio and that their idiosyncratic risk can be described by normally distributed random variables with zero mean and constant variance. Now if the shared market component exhibits large negative jumps, then the two equities are more likely to exhibit large negative returns at the same time than they are to exhibit large positive returns at the same time. The greater in magnitude are the jumps relative to the standard deviation of the market component and the idiosyncratic components, the more extreme the asymmetry becomes. Taken to its extreme — setting the idiosyncratic risk to zero and only focusing on jumps in factors —is effectively the standard approach to stress testing. While this approach may lack the mathematical elegance of copulas and higher order cross moments, it is extremely practical. It is also extremely powerful. For any single factor, we can use all of the tools described in the first section — ARCH, jump diffusion, skewness, kurtosis, etc. — to create an extremely diverse set of joint distributions. When only a single factor is involved, this approach is both simple and transparent.

While all three of the approaches represent a significant improvement over models that rely on correlation alone, it is not clear that one is better than the others. That we have not settled on a standard approach has undoubtedly led to less work being done in this arena, both in academia and in practice.

Future = Past

Risk management has an existential problem. At some level almost all of our models assume that the future is going to look like the past. Our models are backwards looking. Even when our models don't rely on historical data directly, they are likely to do so indirectly. But what if the future is not like the past? While this may seem like a very important question, risk managers likely spend far too much time worrying about this problem.

While this question might seem particularly important to risk managers, it is a problem that all forecasters face. In fact, economists have long referred to the situation where an event that is not present in the historical data, but has a non-zero probability of occurring in the future as the "peso problem". Prior to 1976, the Mexican peso was pegged to the US dollar. During the 1970's, market participants (correctly as it turns out) feared that the peso would be allowed to float against the dollar, but there was no way to quantify this probability based on the historical exchange rate. There were zero observations of the peso being unpegged, but the probability that it would be unpegged was clearly not zero. More recently Nassim Taleb has popularized the term Black Swan to describe the fact that we are likely to underestimate the probability of these possible, but rarely or never seen events.

At some level this question of "how can we be sure that the future will resemble the past?", becomes a question of "how do we know anything?". Modern philosophers from Descartes to Locke to

Hume to Kant have struggled with these questions for hundreds of years. Hume famously struggled with what could be seen as the opposite of the peso problem: just because we have seen the sun rise every day for as long as we can remember, why is it logical to suppose that it will rise tomorrow? Starting in the 19h century the Existentialists looked at this mess, and decided that there were some problems that we just can't answer. It's not that the questions —Do we exist? How do we know? Why are we here? — are not important — they are extremely important — it is just that we do not possess the faculties to answer these questions.

Most of the time, the future is very similar to the past. Using historical data as the basis of our risk models is not a bad idea. Of course there will be times, as with the peso problem, when events that are not present in the historical data do occur; however, these events are infrequent, and often easy to foresee. Still other events are infrequent and impossible to anticipate, no matter how much we try. We simply don't possess the faculties to foresee some events, and we have to accept that.

If there are consequential events that are impossible to foresee, the answer is not to spend more time trying to foresee these events, but to try to make our systems more robust to these unforeseeable risks. This can be done through hedging, or by building reasonable buffers into our risk systems.

This paper has framed today's challenges in risk management using a theory that is nearly 60 years old. In the time since Modern Portfolio Theory was introduced we have made significant progress, and while the future is not entirely predictable, many things are very likely. The sun will likely rise tomorrow and we will likely continue to make progress in the three areas of quantitative risk management outlined above. The future of quantitative risk management is very bright.

References

Engle, Robert (July 1982). “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation”. *Econometrica* 50 (4): 987-1007.

Markowitz, H.M. (March 1952). "Portfolio Selection". *The Journal of Finance* 7 (1): 77–91.

Merton, Robert (1976). “Options Pricing When Underlying Stock Returns are Discontinuous”. The *Journal of Financial Economics* (3): 125-144.

Miller, Michael. 2012. *Mathematics and Statistics for Financial Risk Management*. New York: John Wiley & Sons.

Taleb, Nassim. 2001. *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Random House.