



<http://www.bankinvestmentconsultant.com>

[Back](#)

The New Non-Normal

By [Bryce James](#)

July 1, 2010

Investment risk means many different things to different people. Yet for decades the investment world has relied primarily on one particular measure of risk, standard deviation, which I believe doesn't apply very well to today's markets. Many of the violent swings that we've seen would never have been predictable under traditional standard deviation calculations. Let me show you why and introduce you to another risk measurement concept that may more accurately reflect reality.

[Like what you see? Click here to sign up for our daily newsletter to get the latest on advisor market trends, investment management, retirement planning, practice management, technology, compliance and new product development.](#)

Today's standardized risk metric took hold in the 1950s when top economists from major academic institutions, such as Maurice Kendall (London School of Economics); Paul Samuelson (MIT); Harry Markowitz (University of Chicago); and William Sharpe (UCLA) observed that over the long term, changes in security prices plotted on a histogram chart resembled the shape of the symmetrical bell curve.

The plot of the data in [Figure 1](#) shows that most days prices were barely up or down from the long-term, mean-average price movement, while large changes in price occurred far less frequently. "Long term" generally refers to 40 or more years of daily prices.

Since the data approximately fit a familiar bell curve (that came with well-developed mathematical tools), it was tempting for the early researchers to adopt it as a good fit. The ease of plotting a bell curve made the chore of estimating and interpreting risk easy: The wider the bell-i.e., the wider the variance or average distance from the mean-the greater the risk.

This variance is measured using the standard deviation calculation-the same method your teachers used to score your tests back in school. The assumption is that 68.3% of returns will fall under one standard deviation (1σ), 95.4% of returns will fall under 2σ and 99.7% will occur within 3σ . Low-volatility securities, like short-term bonds, have very low standard deviations, such as 2%, whereas volatile securities, like emerging-market stocks, will have large standard deviations, such as 24%.

In the theoretical world where security returns can be described by a normal bell curve, you might assume that the expected return of a specific security is 8% per year and its standard deviation is 4% (see [Figure 2](#)). Then you would expect your daily returns to be between 4% and 12%, 68.3% of the time; and between 0% and 16%, 95.4% of the time; the returns will fall between -4% and 20% 99.7% of the time. Furthermore, a 1σ event should

occur every six days; a 2σ event once a month, a 3σ event once every 1.5 years, and a 4σ occurrence once every 63 years.

In the 1950s, Harry Markowitz took the concept a step further. Realizing that most investors are concerned with the probability of losing money, he focused on downside risk, by examining just the left side of the frequency distribution. This concept (Figure 3) is called semi-variance because it measures only the half of the variance with losses. With semi-variance the probability that a loss will be more than 1σ is 15.9%; more than 2σ is 2.3%; more than 3σ is 0.14%; and more than 4σ only 0.003%.

In the 1980s, risk modeling really took off, along with the introduction of [Modern Portfolio Theory](#), which uses standard deviation as the foundation of its asset allocation methodology. Much of the impetus came from bank and insurance companies that were trying to manage risk in the wake of the savings & loans crisis. ERISA laws and compliance regulations were also spurring brokerage firms to find better ways to mitigate their liability.

In 1993, two J.P. Morgan analysts realized that banks were less concerned about small losses caused by a house fire, flood or cracked foundation. Their clients were concerned about catastrophic losses. They wanted to know the financial effects of a natural disaster such as a hurricane or an earthquake.

To calculate the odds of such worst-case scenarios, they took the concept of semi-variance and moved the marker toward the far left tail of the distribution. They typically measured the last 1% of the distribution, located at about 2.6. They coined this risk method "VaR" to signify how much "Value" they would have "at Risk" in their portfolios should the worst occur. With VaR (Figure 4), a risk manager would cite a 1% chance of losing x% on a given day.

WRONG TURN

But there is a common limitation to all of these risk methods (standard deviation, semi-variance and VaR). They all assume that distributions are normally distributed. But [distributions tend to be lopsided](#) in one direction or another, a tendency called "skewness" in the world of statistical modeling. Normal distributions and standard deviations describe many, but not all, natural phenomena and they tend to work well in low-volatility markets. But these measurements don't apply to the more recent, complicated behavior of the securities markets. In nature almost anything you can measure can fit into an idea of random variables, which tend to follow a bell curve, the exceptions, like hurricanes, can be applied to commodities, which relate to seasons.

But add investors or newfangled investment vehicles to the mix and you get enough unpredictability to where the bell curve just doesn't fit the data. For example, a five standard deviation natural event should only occur once in 7,000 years, according to traditional theory, when in fact such events occur on average every three to four years in the financial market.

Even though a 22σ market crash occurred in October 1987, traditional risk models would have you believe that a 10σ event should occur once in 260 million trillion years and the 22σ day should happen only once in three lifetimes of the universe! These infinitesimal frequencies aren't remotely close to observed reality. If a month is the second worst in 80 years, then the odds of occurrence is 1 in 40, not once in a trillion years as forecast under a normal distribution.

Why our industry continues to rely on outdated normal distributions is perplexing, yet this method is taught at most schools and credentialing organizations. Major banks calculate their risk using normal standard deviation-based VaR under the requirements of the international Basel II Accord. Financial planning models, insurance illustrations and investment consulting tools illustrate their results with standard deviation, assuming that the worst one-year loss in equities will be 17% and losses of 40% or more should never occur in a lifetime.

A NEW APPROACH

The market crisis of 2008 is yet another example of how these models fail. A normal distribution model ignores the tails of distributions, or the events that occur outside three standard deviations. The October 2008 drop of -16.79% was the ninth worst month in U.S. history, with the all-time worst month of -29.73% having occurred in September of 1931. Converting to non-normal distributions (which can be calculated in several ways) brings the

true probability back in line with reality. Non-normal distributions more accurately measure risk and predict a significant number of extreme events, also known as outliers because they describe events that fall outside of the bell curve distribution. Most events in market history fit within non-normal distributions (see [Figure 5](#)).

[Figure 6](#) shows the non-normal distribution of the Vanguard Emerging Markets ETF (VWO). Note the large number of events that occurred in the left tail. Looking at this market data, a statistician might say that it displays excess kurtosis (outliers) and a negative skew (extended left tail). Translated, volatile days in the market are more frequent than predicted and these volatile days tend to result in extreme gains and losses.

Remember that VaR assumes a normal distribution to calculate the loss that should not be exceeded-some percentage of the time, typically 99%. These loss estimates, which we falsely believe to be true, are frequently low because of the incorrect normal distribution. As a result, what we expect to lose 1% of the time is usually much greater than predicted.

Our approach starts with non-normal data distributions to do the VaR calculation. Then considering how "fat" the left tail of the distribution may be we estimate the size of that loss, which is known as Conditional VaR (CVaR), i.e., how big a loss should be expected in the 1% worst-case scenario. Take for example, your basic 60% equity/40% bond asset mix. Where VaR might indicate you have a 1% chance of losing (at least) 15%, CVaR would estimate that you have a 1% chance of losing 22%-a difference of 7%.

The accuracy of any risk metric can be determined by comparing its forecast losses with its actual results. For example, the VaR calculated at the 99% level for the S&P 500 daily returns from 11/15/2001 through 2/5/2010 (about 2,130 trading days) would have violated its estimated loss 31 times compared with only six times with CVaR. In other words, if you calculated that period you would have lost more than what you estimated 31 times versus only six times beyond your loss limits with CVaR.

To estimate the risk of more extreme losses, the VaR or CVaR setting is raised from 99% to 99.5%, where losses exceeded the VaR level 13 times versus only four with CVaR. So not only did losses exceed the VaR level much more often than expected, but 13 of those losses were large enough to be considered very rare. Investors whose expectations were based on CVaR had fewer unpleasant surprises, i.e., significant losses.

WORST CASES

Imagine you are a billion-dollar hedge fund using leverage of 20-1. Using VaR you expect the most your portfolio stands to lose on a given day is 2%, so you sleep well knowing that you have twice as much capital as you need to absorb losses. Then comes that unpredictable day where the account is down 5% or 3% more than the estimated "worst case"-an occurrence that should have only happened once in a lifetime. A 5% loss with 20 times leverage places the equity value of your billion-dollar hedge fund at less than zero by day's end. Add the complications of mark-to-market accounting and illiquidity, and you get a small taste of what happened to the Nobel laureates who relied on VaR at Long-Term Capital Management in 1998.

Mathematical models are often created for the purpose of replicating nature. Financial risk models based on assumed normal distributions have been proven to be wrong repeatedly, so why continue teaching and using them, costing companies and investors billions if not trillions in unexpected losses?

The study of non-normal distributions is not new. Studies in commodity prices demonstrating non-normality date back to the turn of the last century. Benoit Mandelbrot, co-creator of chaos theory, inventor of geometric fractals and authority in Extreme Value Theory wrote about the weaknesses of normal distributions in finance in 1963, as did Emil Gumbel in 1958. Unfortunately Wall Street seems too dependent on easy-to-understand/explain models.


Changing the assumed shape of the distribution is a meaningful first step to improving risk measurement. Additional improvements can be accomplished by analyzing the data with tools such as GARCH, which is a formula for observing the clustering of data as it moves through a distribution. This concept won the Nobel Prize in 2002.

There is ample research on these progressive methods and plenty of talent from schools offering a Master's of Science in subjects such as Financial Engineering, Computational Finance or Mathematical Finance; just don't look in the traditional finance or economics department to find them.

However, there is hope for change. According to Doug Martin at the University of Washington, the number of schools offering these advanced programs has blossomed from a handful in 2000 to some 20 in the U. S. and 35 internationally, as of 2004. It's estimated that the number doubled again as of January 2010.

I encourage our industry to upgrade their best practices in risk management. This includes upgrading the Basel II risk metrics for the international banking standards. The responsibility is equally shared with our educational associations (that grant certifications) such as PREMIA, GARP, CFA, IMCA and CFP. The proof is there: The old way is broken. It's time to adopt a superior method of measuring risk, starting with CVaR.

Bryce James is president of SmartPortfolios, an RIA in Seattle. Contact him at bjames@smartportfolios.com.



[About Us](#) | [Contact Us](#) | [Advertise with us](#) | [Privacy Policy](#) | [Terms of Use](#) | [Site Map](#)

© 2010 Bank Investment Consultant and SourceMedia, Inc. All rights reserved. SourceMedia is an [Investcorp](#) company. Use, duplication, or sale of this service, or data contained herein, is strictly prohibited.