Revisiting Modern Portfolio Theory and Portfolio Construction

By Bryce James

Years ago, I was speaking at an investment industry conference when a man in the front row yelled out, "Is this just another rubber chicken presentation on asset allocation?" I responded with, "Fasten your seatbelt!" and I'll say it again here, because this is not your generic article on post-modern portfolio theory. I'm going to challenge you to question what you think you know about asset allocation and ask you to keep an open mind to what is possible.

Figure 1: Standard Deviation Probabilities from the Mean Return

Probability of Return at 1, 2, and 3 Standard Deviations

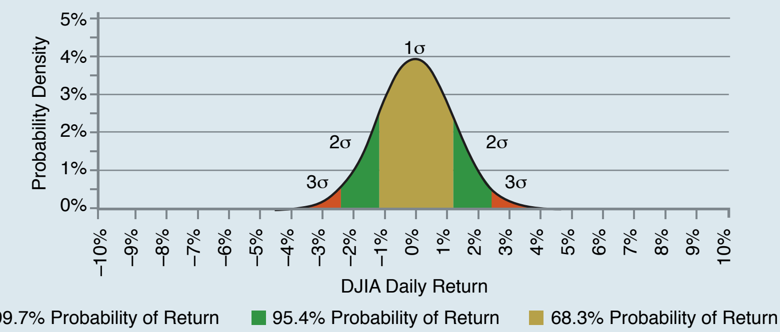

Source: Internal calculation using Excel. Data from Thomson-Reuters

Birth of a New Technology

Computer technology was born in the 1950s with the first commercial computer—the UNIVAC 1—delivered in 1951, followed by the first hard disk in 1955, Fortran computer language in 1957, and the modem in 1958. This era also gave birth to modern portfolio theory (MPT). You know about Harry Markowitz and what he accomplished. All asset allocation work henceforth is a manifestation of his concepts, founded on the principles of diversification. Implementing MPT is to apply a process called mean-variance optimization (MVO).

An MVO model is a three-step process for portfolio construction:

1. Collect a large amount of historical price data on each security in your fund universe (typically 30 years of daily prices and dividends).
2. Calculate the risk and return of each security and the correlation between pairs of securities. In traditional MPT, this is performed using standard deviation (σ) for risk, mean return (μ) for return, and linear correlation (ρ) for correlation.
3. Optimize the portfolio using a covariance matrix to seek the optimal efficient frontier, i.e., asset mix.

In a classic MPT model, which is based upon probability theory, many assumptions are made. The core assumption was founded on the work of Louis Bachelier (Mandelbrot and Hudson 2004, 9), who in 1900 introduced a probability-based model following the random walk theory, which is named after the unpredictable path of a drunken man (Fox 2009, 40–41; Sornette 2002, 38). Bachelier postulated that the market is like a 50-50 coin toss and that most market moves are measurable. This led him to conclude that markets are efficient. He calculated that 68.3 percent of the changes up and down are small, falling within one standard deviation (1σ) from the mean return; that 95.4 percent of the changes fall within 2σ; and that 99.7 percent fall within 3σ. Markowitz (1952, 1959) fit snugly into this random walk theory, which academics named the efficient market hypothesis (EMH). The person credited with this integration of the EMH was Eugene Fama at the University of Chicago.

Standardized risk metrics in finance took hold in the 1950s when top economists, such as Maurice Kendall (London School of Economics), Paul Samuelson (MIT), Harry Markowitz (University of Chicago), and William Sharpe (University of Washington and University of California, Los Angeles), observed that over the long term, changes in security prices plotted on a histogram chart (frequency distribution) resembled the shape of the symmetrical bell curve (Gaussian distribution) (see figure 1). Combining all these pieces gave birth to the "science of investing."

As MPT took hold in the early 1960s, folks such as Jack Treynor and William Sharpe introduced the capital asset pricing model (CAPM), improving upon Markowitz's capital allocation line. Risk finally could be measured, as ratios, in various ways through the metrics created by Sharpe, Treynor, and Sortino. Today, these theories, known for their risk metrics, stand as the foundation of investing. But as technology has significantly changed over the past half century, is it possible finance has been asleep?

MPT assumes all information is priced into the security and that yesterday's price has no influence on today's price; each price change is independent from the previous move. The market, and each security, is fairly priced and moves are completely random, with no regard for human emotion. Therefore, MPT assumes: (1) returns are normally distributed, (2) risk (σ) is accurately captured using a normal distribution, (3) historical prices (mean-variance) can be used to predict future prices, (4) linear correlation accurately represents the relationship between pairs of securities, (5) data acquired from data providers is accurate, and (6) markets are efficient, thus ignoring the behavioral science of fear and greed. I'm here to tell you that all these assumptions are false. Stock markets do not follow random walks (see, e.g., Lo and MacKinlay 1988; Grossman and Stiglitz 1980; Colander et al. 2009).

It's so hard to let go of judgments and beliefs. It's human nature to hang on to what's familiar. In our business, it's best to stay safe by doing what everyone else does. We also stay safe by keeping things simple. If it's not simple, clients' eyes roll back in their heads and you've lost the sale. This is why even the newest financial technology—robo-advisors—rely upon the MPT/MVO math from 1952; it's fast and easy. The first person to expose the naked truth about the mathematics was Benoit Mandelbrot (1964). He criticized the use of normal distributions and warned that markets are influenced by human behavior. He should know—he worked with Markowitz, and he was Fama's PhD advisor.

Asset Allocation Engine—Looking under the Hood

The basic parts to the asset allocation engine include data (including time series), risk, return, correlation, and optimization.

Data

In 2005, a law firm representing a group of investors hired me as an expert witness. The clients, all retired, claimed their broker was imprudent in managing their assets, creating losses exceeding 40 percent. The brokerage firm's asset allocation software, based on an MVO model, recommended the usual 60/40 mix. However, the firm chose to use 18 years of historical data rather than the 28 years recommended by the software vendor. As a result, the Monte Carlo simulations estimated a 16-percent return on equities into the future instead of the 9-percent estimate (as of December 31, 1999) using the software vendor's default 28-year time series. Yet this was not the main issue. The broker himself, using his own personal asset allocation software, decided to load only three years' worth of history (ending December 31, 1999) into the asset allocation software. Recall that this was one of the best-performing periods for the stock market, if not the best, in history. In doing so, the broker's MVO model forecasted a 27-percent return, which he presented to his retirement-age clients. As you know, the dotcom bubble soon burst and it took the S&P a decade to break even and 14 years for the NASDAQ. The case was an easy win for our side. Selection of time series is a form of momentum.

The time series is such a critical piece of data, yet so overlooked. I argue that all models are momentum models, because any time you select a time series, you are making an assumption and making a bet. To buy and hold is making a bet that past performance is indicative of future results, contrary to our marketing disclaimers.

Back in the days of the UNIVAC 1, the term GIGO (garbage in, garbage out) was born. Financial companies have been built around delivering quality data. I'm here to tell you it's still problematic. You've seen the Ibbotson chart for most of your career—the one showing the growth of $1,000 over time for each major asset class. What if I told you that one of the single biggest indicators was grossly wrong for decades? It turns out the creator of the Lehman Brothers Bond Index (now Barclays), Ron Ryan, called out Ibbotson for using callable bonds instead of non-callable bonds. In addition, they were using bond proxies when no long bond was issued (which is most of the time) in calculating the long-term bond index. The net effect was an over-allocation to equities and an undervaluing of bond returns. The net result, according to Ryan, created a massive performance difference of 45.3 percent (from 1941–1991). Furthermore, investors would have experienced a standard deviation of 9.04 percent using the Ibbotson Long Treasury portfolio compared to 5.96 percent standard deviation with the actual Treasury composite; a difference of 45.6 percent more risk (Ryan 1992).

In a response to Ryan's claims, Ibbotson colleagues Laurence B. Siegel and Scott L. Lummer (1993) wrote:

> *Today, however, only a very naïve investor would use our 20-year constant maturity series as a benchmark for evaluating a diversified bond portfolio. Mr. Ryan makes a valid point in suggesting an asset allocator could be fooled by the Ibbotson data into underweighting in bonds. This is a danger only if the asset allocator literally believes the disappointing historical return on long-term bonds will be repeated. Again, this is a naïve view. The expected return on a default-free bond is its yield. This, not the historical return, should be used for asset allocation.*
>
> *Asset allocation models should be expectational; they require expected returns as inputs. For bonds, the expected return is observable in the market as a yield and does not have to be extracted from history. To use a historical average as the expected return for bonds, as Mr. Ryan has done, is conceptually improper and leads to poorly constructed portfolios.*

No apologies were given by Siegel, Lummer, or Ibbotson; they just moved blame to the user. But note that Ibbotson no longer shows long bonds on its charts.

Risk

What is risk? If we agree that risk is a measure of volatility (thereby agreeing to ignore all other risks), then it's prudent to compare risk models. We have two classes of risk models, one based on normal distributions and another on a class of non-normal distributions. Normal distributions include the risk metrics: standard deviation (entire variance), semi-variance (downside variance), and Value-at-Risk (VaR), or tail risk (see figure 2). Normal distributions are measured following a Gaussian bell curve.

Table 1 shows the odds of exceeding each standard deviation. For example, in the EMH, the odds of a 4σ event is every 123

Table 1: Odds of Exceeding Each Standard Deviation

Risk	MPT Estimate	Actual	Standard Deviation Error
1σ	6 days	8 days	18% less
2σ	44 days	35 days	27% more
3σ	3 years	0.54 years	5.4x
4σ	123 years	1.9 years	65.2x
5σ	14,000 years	3.6 years	3,828x
6σ	4 million years	10 years	~4,000x
7σ	3.052 billion years	15 years	~201 million x
8σ	6.279 trillion years	30 years	~208 billion x
9σ	34.611 quadrillion years	30 years	~1143 trillion x
10σ	513 sextillion years	60.5 years	~8.5 sextillion x

Figure 2: Three Measures of Risk (Normal Distributions)

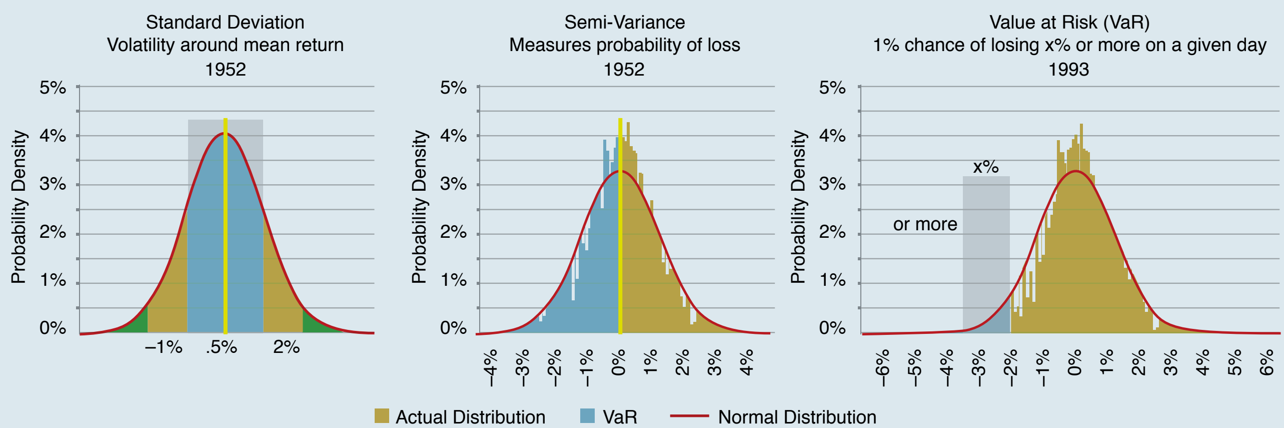

Source: Internal calculation using Excel. Data from Thomson-Reuters

Figure 3: DJIA Daily Returns—Outliers

Normal, Actual, and "Heavy Tail" Distribution (left tail) 1985–2016

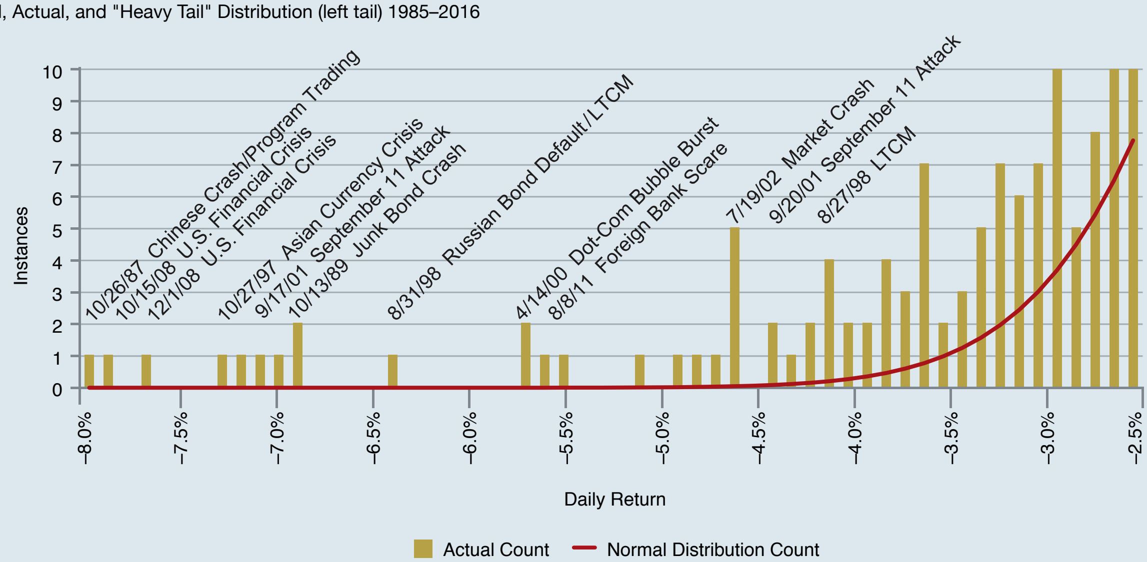

Source: Internal calculation using Excel. Data from Thomson-Reuters

years, when in actuality it's every 1.9 years. In EMH, the odds of a 2008 event is one in three lifetimes of the universe, not every 60.5 years, which is what has occurred historically. That is why we have so many outliers, because of bad mathematics.

VaR moves beyond the probability of losing money (semi-variance) and aims to measure catastrophic risk, known as tail risk. VaR picks a point on the distribution, say the last 1 percent of the bell-curve area, and estimates the probability of losing more than the average loss of the tail surface area. For example, SPY has a 1-percent chance of losing −2.09 percent or more today. VaR grew so popular that it became the standard risk measure for all banks and insurance companies as part of the Basel II Accord. But how does one measure tail risk when one can't see the tail? So far within my research, we have yet to find an equity security that fits within the paradigm of a normal distribution. VaR definitely did not help Long-Term Capital Management (LTCM) in 1998, nor the banks and insurance companies during the 1987, 2000, and 2008 crashes (see figure 3).

VaR's core problem is that it uses an assumption of a normal distribution. Normal distribution can measure only two moments: mean return and variance (how far a distribution is spread out from its mean). You can measure for normality using the Jarque-Bera normality test, which is a goodness-of-fit test of whether sample data have the skewness and kurtosis matching a normal distribution. But actual return distributions have third and fourth moments, called skewness, which is an expression of how lopsided a distribution appears, and kurtosis, which estimates the extent of fat tails in the distribution of data. The red line in figure 3 depicts the probability distribution of a normal distribution. Any gold bar above the red line represents an outlier. Table 2 lists specific daily losses, the date, and the probability of occurrence using the bell curve. At what point do we raise the white flag and surrender the normal distribution?

Non-normal distributions (NNDs) allow more than two moments to be seen. A wide selection of NNDs can be fitted such as Student's-*t*, log-normal, skewed-*t,* paired-*t,* stable, and combinations of distributions. Any of these NNDs can significantly improve your results. For perfectionists, scientists have nice mathematical tests to compare and contrast forms of distributions.

In addition, you can find which method has a history of producing the least number of outliers (aka exceedances, extreme events, black swans); this can be measured at both tails, in absolute numbers and as a percentage. Now you have more-accurate risk tools replacing the inputs of the old ratios of Sharpe, Sortino, and Treynor. Beyond VaR, better tail metrics, such as Conditional VaR (CVaR), can now be better applied.

Utilizing an NND significantly improves your estimation of loss and lowers outliers with CVaR. However, you still are using a single long-term average number to predict the future risk (and return). Fortunately there are models that allow for more-recent information to gain more importance, such as exponentially weighted moving average (EWMA) and GARCH (generalized autoregressive conditional heteroskedasticity). The men who created GARCH, Clive Granger and Robert Engle, won the Nobel Memorial Prize in Economics in 2002.

The magic of GARCH is its ability to capitalize on the newer information, like EWMA, but discount the data as it moves too far from its traditional mean-return. The disadvantage of GARCH is the whipsawing effect during trading ranges in the market. Our research indicates the positives far exceed the negatives with GARCH.

Figure 4 shows SPY, the exchange-traded fund (ETF) replicating the S&P 500. The green line is the price growth of SPY; the black line is the daily price change (up and down); and the red line is a rolling three-month VaR. Note how frequently the black line exceeds the red line to the up and downside. Focusing on the downside, the number of exceedances over the past two years (ending November 4, 2016) at the 1-percent VaR level totals 17. The blue line is a combination of an NND and GARCH; we call this expected shortfall with GARCH features, or ES for short. Using ES, for the past two years, we experienced three exceedances at the 1-percent level versus 17 with a rolling VaR.

Nobel Prize-winning physicist Philip Anderson states:

> *Much of the real world is controlled as much by the "tails" of distributions as by means or averages: by the exceptional, not the mean; by the catastrophe, not the steady drip; by the very rich, not the "middle class." We need to free ourselves from "average" thinking* (Ramalingam 2014, 219).

Table 2: Seven Worst Days (in 30 Years) and the Probability of Loss (Gaussian Bell Curve)

Return	Date	Description	Probability (Years)
–22.61%	10/19/1987	'87–88 Bear Market, Black Monday	377,928,949,357,521 + 25 more zeros
–8.04%	10/26/1987	'87–88 Bear Market, Program Trading	22,546,897,547
–7.87%	10/15/2008	Financial Crisis Bear Market	7,348,618,460
–7.70%	12/1/2008	Financial Crisis Bear Market	2,449,539,487
–7.33%	10/9/2008	Financial Crisis Bear Market	241,966,705
–7.18%	10/27/1997	Asian Currency Crash	97,981,579
–7.13%	9/17/2001	9/11/2001 Attack	72,811,999

I've never met the person who can accurately predict stock prices. So the next question is: Can you predict volatility? Volatility is most notably tracked with an instrument called the VIX. The VIX is the ticker symbol for the Chicago Board Options Exchange (CBOE) Volatility Index, which shows the market's expectation of 30-day volatility. Amazingly, ES with GARCH features is highly correlated (as an inverse) to the VIX, which is based upon the Black-Scholes pricing model. GARCH is using historical distributions of daily prices. In comparing the VIX (S&P 500) to SPY (ETF proxy for S&P 500) we find a very high correlation (see figure 5).

The advantage of ES with GARCH features is the ability to create a VIX-like indicator for every security or portfolio and also perform a fairly accurate forecast of risk. Figure 5 illustrates the similarities between the VIX and ES using the ETF: SPY.

Essentially, ES with GARCH can be a predictor of price. In other words, you can predict price by forecasting risk, what I call "directional risk," and by analyzing the level of risk, what I call "risk regime." Figure 6 shows directional risk in relation to price. Magenta lines (on the upper chart) are decreasing prices, and on the lower chart increasing risk; visa-versa for the purple lines.

Note the high correlation between price and volatility. Note also that fractal behavior and the power law come into play during the 2008 crash, just as Mandelbrot and Hudson (2004) wrote in *The (Mis) Behavior of Markets*.

Return

Recall that Siegel and Lummer (1993) said that "only a fool" would rely on historical returns. Yet they taught it to all of us for decades. The truth is, the entire academic and professional organizational network taught us MPT and its derivative methodologies, such as CAPM, arbitrage-pricing theory, and the three-factor models.

As we reflect upon mean-variance for estimating returns, we must accept that it ignores market cycles, the momentum

Figure 4: SPY—State Street Global Advisors (SSgA) S&P 500 Index

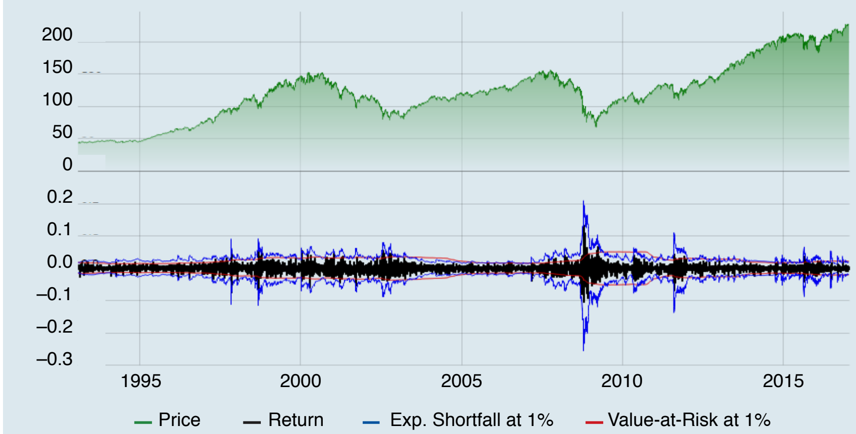

Source: Internal calculation using Excel. Data from Thomson-Reuters

Figure 5: Dynamic Risk (S&P, SPY)

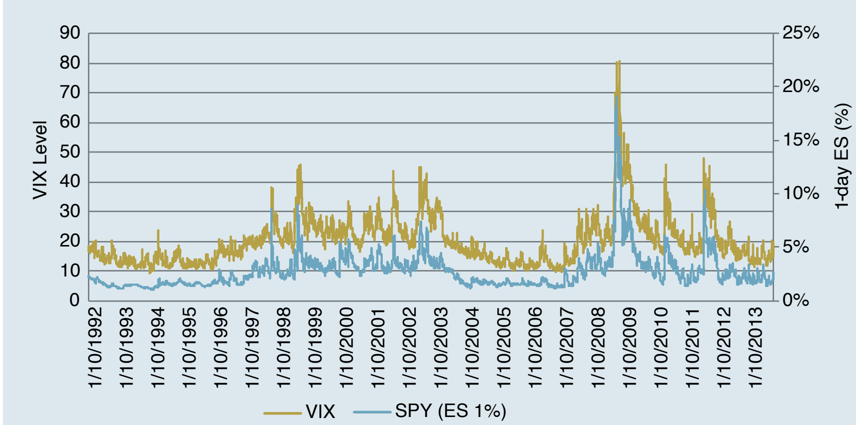

Source: Internal calculation using Excel. Data from Thomson-Reuters

Figure 6: SPY—State Street Global Advisors (SSgA) S&P 500 Index

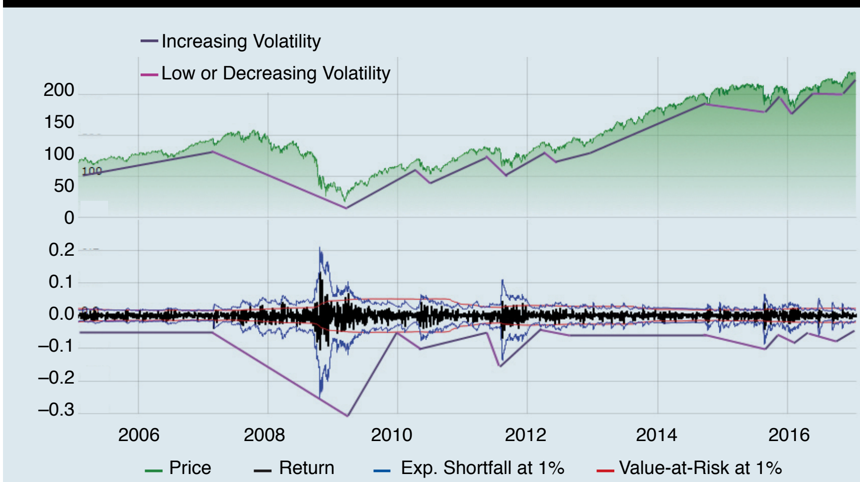

Source: Internal calculation using Excel. Data from Thomson-Reuters

Figure 7: Correlation between Two Securities

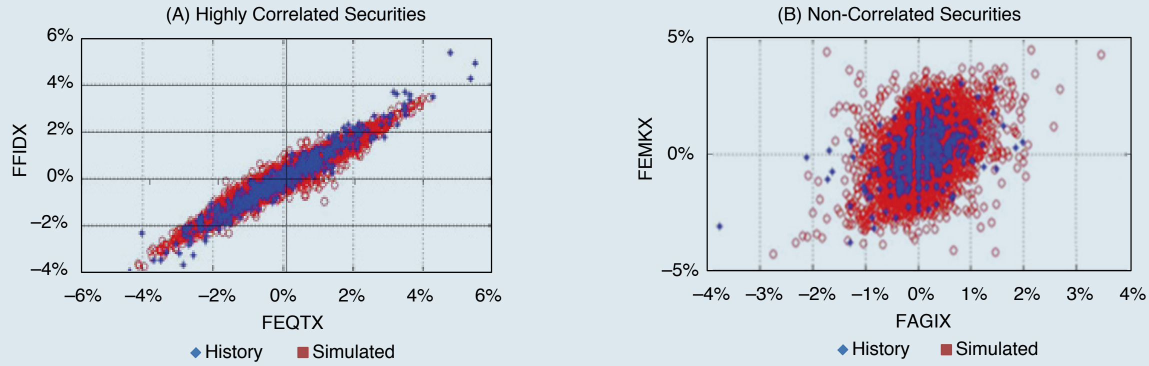

Source: Internal calculation using Excel. Data from Thomson-Reuters

effect, economic and geopolitical events, and social and technological trends. We must accept that Brownian motion might work well in some sciences, such as botany, but not when emotions or dynamic changes in market forces come into play. The world is more fractal than linear, which is why we replaced the *Farmer's Almanac* with Doppler radar for forecasting the weather. Forecasting returns using Monte Carlo simulations with MPT/MVO is the *Farmer's Almanac* of finance.

Forecasting returns is as problematic as the weather. Can you really count on analysts to give you accurate forecasts? Are they not biased by investment banking fees or group think? So do we use historical mean-variance, or an analyst's estimates, or third-party research, or create our own, and then apply a tool like the Black-Litterman model (BLM)?[1] In a way, BLM is not so different from GARCH models. BLM uses long-term performance as a baseline for estimating returns, then brings in views, such as estimates (with probabilities), to tilt the return forecast. Using robust research and testing methodologies, one could create a superior return function using one or more views. Each view can be whatever you want it to be, such as a technical strategy, a trading algorithm, or a collection of fundamental or economic data. There is no way for me to compare all possible return functions here. What I can say is that mean-variance is way down the list of optimal return functions. After all, past performance is not indicative of future results.

Correlation: The Golden Goose

Correlation is the foundation of asset allocation because diversification is the underlying principal for reducing risk and enhancing returns. The importance of reducing nonsystemic risk is the cornerstone of Markowitz's work, but it's really driven home by Brinson, Hood, and Beebower (1986) and Brinson, Singer, and Beebower (1991). Xiong et al. (2010) completely re-analyzed the situation and made a stronger case for where returns are sourced.

Xiong et al. (2010) reveals that the market movement component accounts for about 80 percent of the total return variations of a portfolio. Of the returns in excess of market return (i.e., attributable to manager performance), 90 percent of variability of returns across time is explained by asset allocation policy, 40 percent of variation between funds is explained by differences in asset allocation policy, and 100 percent of return amount is explained by asset allocation policy. Furthermore, with market movements removed, asset allocation and active management are equally important in determining portfolio return differences within a peer group, with each accounting for around 20 percent.

Correlation is still a key driver to returns and risk. However, like distributions, not all correlations are the same and linear correlation can be the most limiting. As intelligent beings, we somehow find it okay to take a bunch of historical plots and draw a least-squared line through them and come up with a single number to represent the relationship between two securities. This is like monitoring a couple who've been married 30 years and averaging their emotions, then defining their relationship as one emotion. Life doesn't work that way. We all know correlations move toward 1 on big up and big down days. Clearly one number cannot represent all the dots all the time, which range from highly correlated to highly negatively correlated (see figure 7).

Copulas are another form of correlation model.[2] By switching to a rank-correlation methodology, such as Kendall's tau rank correlation coefficient, we can explore more-flexible correlation models, including copulas with NNDs.[3] Copulas, in the wrong hands, can be bad. They didn't work well for the banks and insurance companies who blamed copulas for their large losses during the financial crash. A copula is only as good as its underlying distribution. If you run a copula on a normal distribution you ignore the tails, just as VaR, semi-variance, and standard deviation do. Copulas also need NNDs to see the tails.

In figure 8B, note the high correlation in the tails on big up and down days; linear correlation (figure 8A) ignores this dynamic nature. Traditional models such as MPT depend upon linear correlation matrixes to define intermarket relationships. We know correlation, price, and risk are dynamic, not static. Thus the adage that the only thing to go up in a down market is correlation. Likewise, Gaussian copulas do not account for tail correlation.

One such model that does account for tail correlation is Student's *t*-copula. This now means we can model increased correlations during market extremes.

What concerns me about correlation is the long-term trend. Correlations have continued to increase since 2000. Maybe the markets are becoming more efficient, or maybe there is a shortage of traders—who knows? At this pace, everything we know about asset allocation and diversification may go away.

Optimization

Figure 9 shows actual returns for a moderate-risk target date fund offered by Morningstar to establish a baseline risk level. We concluded that the average expected shortfall of this fund over a 10-year period (August 2004–August 2013) was −1.5-percent ES at the 1-percent level. In other words, this fund should lose −1.5 percent or more, 1 percent of the time. The actual results are illustrated as orange lines in figure 9A. By simply converting to ES with GARCH, the exceedances were reduced and performance improved.

The new target-risk methodology (ES with GARCH) still deviates based on intramonth risk changes, but overall these deviations are shorter and smaller in magnitude. Intermonth hedging can also add value (see table 3).

Figure 8: Comparing Dependency Models

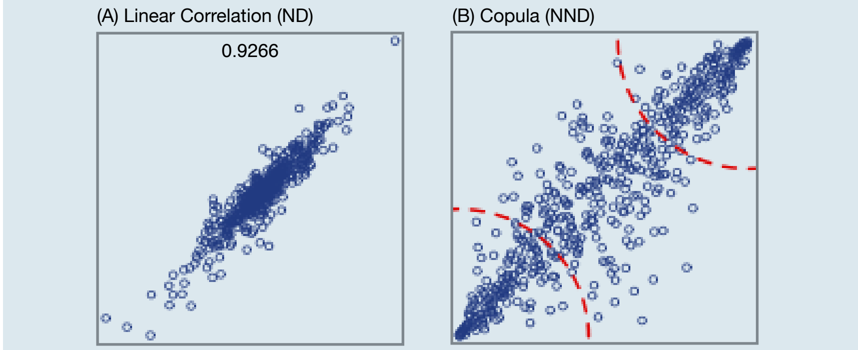

Source: Internal calculation using Excel. Data from Thomson-Reuters

Clearly risk models are not alike; nor are momentum models. Momentum can be achieved with technical indicators, factor models (such as Fama or Black-Litterman), or econometric models. To say that momentum doesn't work is simply not true. Eugene Fama was the father of the efficient markets hypothesis; he built an empire at Dimensional Fund Advisors with his three-factor model. Then in 1997, he read a research abstract by Mark Carhart, who added a fourth factor—momentum—to the three-factor model by using 11 months of previous momentum. Fama quickly added this strategy to his offering and it's been a meaningful contributor to his performance (Bello 2008; Hou et al. 2015). Fama has long since reversed his views of an efficient market.

An Evolution Waiting for You

As we moved through time, some fixes to MPT were Band-Aids and others were evolutionary processes. From MPT came MVO, then CAPM, followed by arbitrage pricing theory with its factor models, then, the Black-Litterman multi-factor model and the French-Fama factor models. It's an evolutionary process, but we remain stuck in the old paradigm because too much

Table 3: Target-Risk Methodology

Target Risk >	New (ES)	Old (σ)	S&P 500
Annual Growth	7.41%	6.16%	7.67%
Annual Volatility	9.33%	10.15%	14.28%
Maximum Drawdown	–25.24%	–36.10%	–50.94%

Figure 9: Realized Expected Shortfall versus Targeted Risk

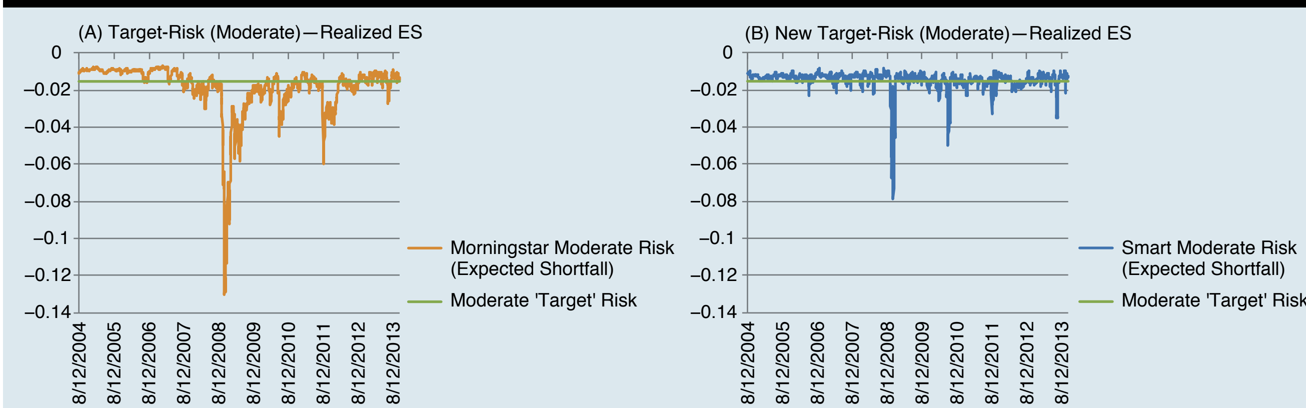

Source: Internal calculation using Excel. Data from Thomson-Reuters

would need to change and some firms could lose significant market share. It's all about the Benjamins. With dynamic models, style boxes would be rendered meaningless, as would investment-policy questionnaires the way they are written today. Case law would need to change as would the onslaught of robo-advisors and consulting firms. Not until more people get sued for outdated methodologies will things start to change.

Imagine you could combine the best of these tools. Imagine that you had access to clean data, that you can choose any data distribution you desire based on valid testing, that you could apply GARCH features to make the risk, return, and correlation numbers behave more like a Doppler radar, that you could modify your return function using Black-Litterman (for the return function) and you could incorporate copulas using NNDs to dynamically calculate correlation. The concept of creating models that scale is the foundation of extreme value theory (EVT). When I combine EVT with GARCH features and newer correlation methodologies, I create a dynamic risk theory implemented with a dynamic optimization model.

Validating Galileo Galilei

MPT, unlike the sun, is not the center of the universe. In 1992 the Vatican, after 350 years, finally renounced condemning Galileo for stating the Earth revolved around the sun. Likewise, markets do not follow the efficient market hypothesis (i.e., the random walk theory). Nor is MPT, MVO, or any risk, return, or correlation model built on normal distributions legitimate in portfolio construction.

Fama (1991) rejected the random walk theory and promoted the idea that expected returns vary with time. Through decades of research and managing assets, Fama has evolved; it's time we all do.

Perhaps Fischer Black (1986) stated it best: "In the end, a theory is accepted not because it is confirmed by the conventional empirical tests, but because researchers persuade one another that the theory is correct and relevant." How long must we keep convincing ourselves?

In the beginning, one man pointed out the flaws in MPT and EMH—Benoit Mandelbrot, a brilliant mathematician whose work advanced myriad fields of study, from physics to finance. In 2004, Mandelbrot wrote a book about finance called *The (Mis)Behavior of Markets*. In it, he stated, "If there is one message I'd like to pass on … it is this: Finance must abandon its bad habits and adopt a scientific method." He also wrote, "Extreme Value Theory, borrowed from the insurance industry, is on the right track; it assumes prices vary wildly, with fat-tails that scale."

It's time to take a hard look in the mirror and ask ourselves if we use the traditional tools because they get the results we want and need or because they're easy and safe and everyone else uses them. We tell ourselves that our planet is at the center of the universe because we fear persecution, but we do so at the risk of being foolish. Watching from the outside is like visualizing Einstein's definition of insanity. I'm guessing Galileo and Mandelbrot might have had a lot in common.

Bryce James is president of Smart Portfolios, LLC and president of Portfolio Mason, LLC. He earned a BS in accounting, finance,and marketing from Central Washington University. He is a former CIMA® certificant. Contact him at bjames@smartportfolios.com.

Endnotes

1. The Black-Litterman model is an asset allocation model that was developed by Fischer Black and Robert Litterman of Goldman Sachs. It is essentially a combination of the two main theories of MPT: CAPM and MVO. The main benefit of the Black-Litterman model is that it allows the portfolio manager to use it as a tool for producing a set of expected returns within the MVO framework. This can allow the manager to avoid certain problems or issues inherent in the MVO framework, such as the concentration of portfolio assets in only a handful of the assets under optimization.
2. In probability theory and statistics, a copula is a multivariate probability distribution for which the marginal probability distribution of each variable is uniform. Copulas are used to describe the dependence between random variables.
4. In statistics, the Kendall rank correlation coefficient, commonly referred to as Kendall's tau coefficient (after the Greek letter τ), is a statistic used to measure the ordinal association between two measured quantities.

References

Bello, Zakri. 2008. A Statistical Comparison of the CAPM to the Fama-French Three Factor Model and the Carhart's Model. *Global Journal of Finance and Banking* 2, no. 2: 1–23.

Black, Fischer. 1986. Noise. *Journal of Finance* 41, no. 3 (July): 529–543.

Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. 1986. Determinants of Portfolio Performance. *Financial Analysts Journal* 42, no. 4 (July/August): 39–44.

Brinson, Gary P., Brian D. Singer, and Gilbert L. Beebower. 1991. Determinants of Portfolio Performance II: An Update. *Financial Analysts Journal* 47, no. 3 (May/June): 40–48.

Carhart, Mark. 1997. On Persistence in Mutual Fund Performance. *Journal of Finance* 52, no. 1 (March): 57–82.

Colander, David, Hans Föllmer, Armin Haas, Michael Goldberg, Katarina Juselius, Alan Kirman, Thomas Lux, and Brigitte Sloth. 2009. The Financial Crisis and the Systemic Failure of Academic Economics. Kiel Institute for the World Economy Working Paper 1489 (February).

Fama, Eugene. 1991. Efficient Capital Markets: II. *Journal of Finance* 46, no. 5 (December): 1,575–1,617.

Fox, Justin. 2009. *The Myth of the Rational Market: A History of Risk, Reward, and Delusion on Wall Street*. New York: HarperCollins Publishers.

Grossman, S., and J. Stiglitz. 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review* 70, no. 3 (June): 393–408.

Hou, Kewei, Chen Xue, and Lu Zhang. 2015. Digesting Anomalies: An Investment Approach. *Review of Financial Studies* 28, no. 3 (March): 650–705.

Lo, Andrew, and A. Craig MacKinlay. 1988. Stock Markets Do Not Follow Random Walks. *Review of Financial Studies* 1, no. 1: 41–66.

Mandelbrot, Benoit. 1964. Random Walks, Fire Damage Amount and Other Paretian Risk Phenomena. *Operations Research* 12, no. 4 (July/August): 582–585.

Mandelbrot, Benoit, and Richard L. Hudson. 2004. *The (Mis)Behavior of Markets: A Fractal View of Financial Turbulence*. New York: Basic Books.

Markowitz, Harry. 1952. Portfolio Selection. *Journal of Finance* 7, no. 1 (March): 77–91.

——. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley.

Ramalingam, Ben. 2014. The Devil Is in the Dynamics. Chapter 11 in *Aid on the Edge of Chaos: Rethinking International Cooperation in a Complex World*. New York: Oxford University Press.

Ryan, Ron. 1992. Bond Data Problems Eyed. *Pensions & Investments* (December 7).

Sharpe, William F. 1964. Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *Journal of Finance* 19, no. 3 (September): 425–442.

Siegel, Laurence B., and Scott L. Lummer. 1993. Criticisms of Ibbotson Bond Data Answered. *Pensions & Investments* (January 11).

Sornette, Didier. 2002. *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton, NJ: Princeton University Press.

Xiong, James X., Roger G. Ibbotson, Thomas M. Idzorek, and Peng Chen. 2010. The Equal Importance of Asset Allocation and Active Management. *Financial Analysts Journal* 66, no. 2 (March/April): 22–30.