Modeling the Future: The Full Monte, the Latin Hypercube and Other Curiosities

by Glenn Kautt, CFP, and Fred Wieland, Ph.D.

Financial planners are investigating the usefulness and accuracy of a new method called stochastic modeling for evaluating alternative planning scenarios for their clients. Most of the questions about stochastic modeling have been focused on the validity and accuracy of the financial and historic data used to drive processing. While these are significant issues, the authors have uncovered others not previously addressed in this industry that may materially affect model outcomes. This article focuses on sampling technique and parameter estimation.

Mr. Kautt is president of The Monitor Group Inc., a financial and investment advisory firm in Fairfax, Virginia.

Dr. Wieland is a scientist for a McLean, Virginia, research company.

A client is sitting with her financial planner, who is patiently explaining a different kind of retirement analysis: “Mary, this chart shows you will have plenty of money during retirement. Look—the graph shows you will have a 90 percent probability of having money even when you are 95 years old!”

Are there planners who really say such things? Sure, there are—and they do so with complete confidence. Whenever a new technology, system or technique comes along in a professional field, there are those who actively embrace it without a second thought. Mary’s planner is one of these people, and she’s recently discovered the wonder of stochastic modeling.

Of course, not everyone is like this planner. A much larger part of the profession cautiously enters into a partnership with new ideas only after much consideration and study. Nonetheless, the rewards of even cautiously adopting these techniques can be an immediate and significant improvement in the information clients have to make decisions about their future. This article addresses the majority of our profession who are still studying the new financial planning techniques of stochastic modeling and want to better understand its strengths and weaknesses as they employ its considerable power.

One of the weaknesses of any model is dependence on accurate inputs. If there are inaccuracies, then the output is inaccurate. Remember the old saw about computers: garbage in, garbage out? Critics of stochastic models say the inputs often cannot be determined with enough precision to make a forecast realistic. These planners argue that inaccuracies can stem from incorrect assumptions about the distributions and averages of future returns, correlations between different investment behaviors, and variances in the client’s financial behavior (read that as changing cash in- and outflow patterns, which radically affect the financial model’s outcome).

Proponents of stochastic models believe that including factors for risk make the outcomes more realistic. Models that do not factor in risk, they claim, ignore the biggest reason for modeling in the first place: the uncertainty of the future. They argue that an assessment of the likelihood of an outcome, even if sometimes inaccurate, is better than a single answer that has a high probability of being wrong.

Where’s the truth in all this?

The critics and proponents of stochastic modeling are really discussing several related issues: parameter estimation (also known as input validity), model design, human behavioral factors and computational validity. Put another way, how accurate do assumptions, data inputs, modeling strategy and computational techniques have to be for realistically accurate answers to result? This article addresses computational and sampling techniques and leaves the other issues to future commentary.

On a Clear Day, I Can See Forever...Right?

Assume we break down the accuracy of stochastic modeling into two basic questions: Is the model sampling the data correctly? Does the outcome accurately indicate the true probabilities of the event?

A simplified comparison can illustrate sampling theory and make it easier to understand. Consider an archer shooting an arrow at a target. Assume you know from past personal observation that the archer hits the target about 75 percent of the time. When he misses the target, you've observed he misses by an average of ten yards. Yet conditions have changed slightly: The archer is a little older, you have reason to suspect his aim is not what it used to be and the weather on this particular day has a strong cross-wind that might make his arrows wander more than usual. Your objective is to figure out by how many yards the archer will miss when he does so under these new conditions.

The analogy to investing should be obvious: We have data on returns and volatility from past years. The investment climate has changed slightly, and we're not sure how much the change actually affects results, but we're pretty sure that the past can help us estimate results in the future.
Returning to the archer, you figure he'll hit the target at the same rate, about 75 percent of the time. So he begins shooting: The first ten shots hit the target. If you stop here, you might falsely conclude the archer will hit the target 100 percent of the time, his aim has improved with age and the weather has no effect on the result. This illustrates the danger of too small a sample.

So you urge him to keep shooting. After 20 shots, he has missed the target only once, and then just barely. If you stopped there, you'd conclude he will miss the target five percent of the time by less than a yard. But he keeps shooting. After 100 shots, you find he has missed the target about 15 times, most misses were within ten yards, but two of those misses caught a strong wind and wound up 40 yards from the target. Interesting. Those two errant shots now change your whole calculation: You now conclude he will miss 15 percent of the time and that the average miss is about 15 yards (averaging the misses within 10 yards with the misses at 40 yards).

So how many shots does the archer have to take before you can get an accurate reading on his ability? This is the essential question for stochastic modeling and the heart of this issue. When sampling from a distribution (the archer's shots, in this example), how many samples should we take? It depends. If you are looking for average behavior around a target, then a relatively few samples are appropriate. If you are looking for outlying statistics, then many more samples are necessary.

In other words, if all you care about is approximately how many times an archer hits the target and how many times he misses, then you are looking at a mean, or average, statistic. If you're concerned about by how many yards he misses when he misses, then you are looking at another statistic entirely. You are looking at the subset of times the archer misses the target, and with that subset, you are computing the average distance missed. This is called an outlying statistic.

In financial planning, we are sometimes concerned with average statistics, and at other times we are concerned with outlying statistics. An average statistic would be how much a portfolio is most likely to be worth some number of years from now, given assumptions about returns, volatility, spending habits and so forth. Another average statistic would be an estimate of the volatility around the mean return.

But suppose your client wants to know the probability of bankruptcy, or the probability of a portfolio growing to be worth twice the mean return. These can be significantly influenced by outlying return behavior. They correspond to our measurement of how much the archer misses the mark. Sometimes the archer will miss to the left (by analogy, a low return); sometimes he will miss to the right (a high return). To get an accurate estimate of the misses, or of the outlying statistics, more samples are needed from the distribution.

Exactly how many samples are needed? In the archer's case, it depends on how inaccurate he is. The answer may be counterintuitive—the more our archer misses the mark, the more samples we need. Put another way, the higher the volatility, the more samples that are needed to assess the outlying statistics. If the archer rarely misses, we can get an accurate fix on how rare the miss is and how much he misses with only a few samples. If the archer misses frequently, then we need more samples to determine by how much he misses. In either case, with only a few samples, we can estimate approximately what proportion of the time he hits (a mean statistic). To estimate by how much he misses (an outlying statistic), we need more samples when he misses more frequently.

Honey, I'm Not Lost, Just a Little Disoriented!

There is a little-known procedure that can help reduce the number of samples when outlying statistics are required. In normal sampling, we draw a random number and map it to a probability distribution as shown in Figure 1. This normal sampling is part of a Monte Carlo simulation. So what's the problem? With a wide distribution based on a large standard deviation and an inadequate number of random samples, we might not have an accurate estimate of "by how much he misses." This is the crux of this issue: Has the data been sampled and simulated correctly so the results accurately indicate the true probabilities of the future event?

Figure 1, a normal distribution frequency plot, shows Monte Carlo sampling with eight samples. The height of the distribution shows the probability a sample will be from that region. The first eight samples naturally cluster around the mean, which in a normal distribution is the highest point. The outlying areas of the distribution—those areas to the far left or right—have not yet been sampled. It will take quite a few samples before one appears in either of the two outlying areas.

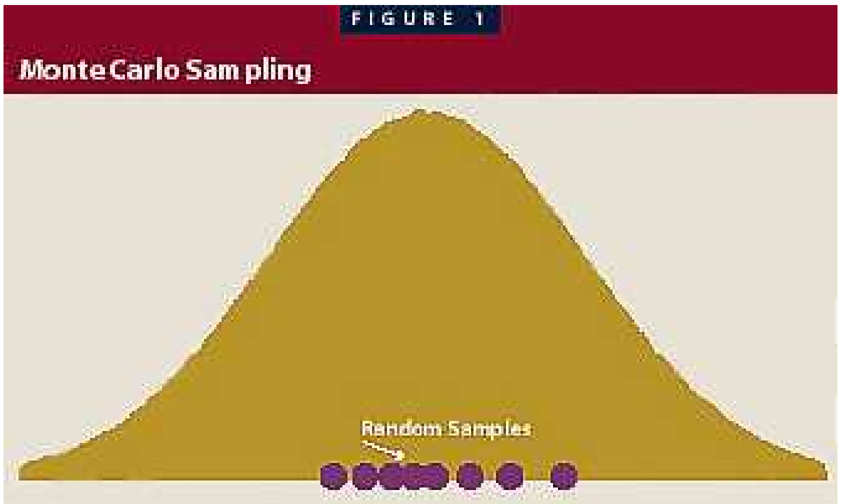

The archer analogy breaks down a little at this point, but with some imagination we can see what is going on. If you know the archer is going to miss 30 percent of the time and you want to measure by how much he misses that 30 percent (the outlying statistic), you can use the following procedure. Tell the archer to take ten shots so that the first seven hit the target and the next three miss somewhere. Continue this procedure for the next ten shots, and so forth, until you have accumulated enough information to compute your result. Of course, our archer has to be cooperative. When dealing with humans, this is a problem, but with computer sampling this is not an issue.

If we know the distribution has 30 percent outliers, then we can draw random numbers such that 30 percent of the time the number will be an outlier. We then use the probability distribution to tell us where the 30 percent of the outliers are. This procedure is known as Latin Hypercube sampling.

Because our archer analogy is simplified, it does not adequately explain the details of Latin Hypercube sampling. The hypercube procedure divides the probability distribution into areas of equal probability (for the mathematically inclined, areas of equal probability density). For example, it divides the distribution into 100 equal areas where each area has 1/100 of a chance of getting an arrow. Here, the hypercube sampling technique criteria assures each area gets an average of one arrow every 100 shots. The procedure then samples each of these 100 areas as if they were different distributions to determine where within each area the arrow is likely to fall. The resulting statistics resemble the original distribution, except the Latin Hypercube technique assures we have more fully sampled the area under the curve with fewer trials.

Figure 2 shows Latin Hypercube sampling. The distribution has been divided into eight areas of equal probability. The bars around the mean are closer together than those at the extremes. Why? The area under the curve between each bar is the same. Eight samples are then drawn—one from each of these eight regions. The samples represent a random point within each region—so the samples can lie anywhere within the region. This procedure generates outlying statistics faster than in normal sampling as the sample technique guarantees points in the extreme regions for each eight-sample set. Remember, in a computer simulation, Latin Hypercube sampling splits the distribution into thousands of regions of equal probability. We have used an example of eight regions here for ease of illustration.

FIGURE 2

Latin Hypercube Sampling

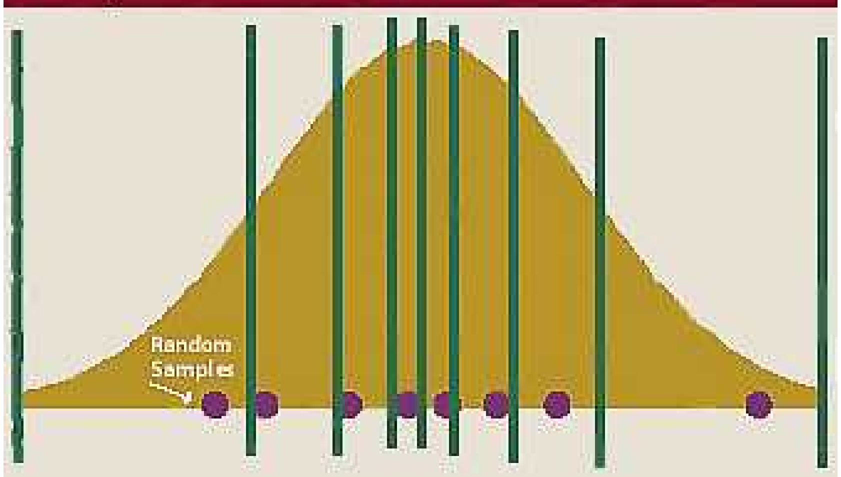

Both normal sampling and hypercube sampling "converge" to the same values when the sample size is large. Using our analogy, you might let the archer shoot millions of times to get the same answer from both procedures. To be technically correct, for smaller samples the Latin Hypercube will be slightly more accurate with outlying statistics and perhaps slightly less accurate for mean statistics.

I Was Lost But Now I'm Found!

Depending upon the variance of the distribution found in investment portfolios, it may take from 100,000 to 1,000,000 iterations to estimate outlying statistics accurately. The number of samples required is higher when volatility is higher and lower when volatility is lower. One way to increase accuracy with smaller samples is to use a Latin Hypercube sampling technique, especially when computational limits exist.

We remind the reader that the number of iterations to reach a statistically valid conclusion depends on all the major factors: sampling technique, input parameters, model design and human behavior. In truth, the impact of these other factors may overshadow the impact of sampling technique. To find out what affects the outcome we must analyze the sensitivity of the model we are using. Before we attempt to do this, however, we must learn a little bit more about the other factors.

Financial planners and investment advisors are justifiably concerned about significant issues in simulation modeling. We now know that with enough samples and by using the proper sampling technique, the future behavior *can* be accurately forecast for both mean returns and their outlying statistics. Yet, because none of the currently available planning models use Latin Hypercube sampling, under certain circumstances a model may not have large enough samples to assure statistical validity given the input data.[1] In a previous article by one of the authors, none of the stochastic planning programs examined had built-in statistical validity tests.[2] As a result, planners who use stochastic simulation planning models must realize the results may produce outlying data that have not been tested for statistical validity.

'Houston, We Have a Problem'

For those of us old enough to remember (with a mixture of amazement and awe), the Apollo moon missions provided us with a sense of accomplishing the impossible. I remember a hot summer night in July 1969 watching a black-and-white TV in Chicago as Neil Armstrong hopped down onto the moon. During the Apollo flights, millions of viewers came to understand how much time and effort went into making sure the trips from Cape Kennedy (now Cape

Canaveral) to the moon and back were perfect, with no problems. A short time later, the flight of Apollo 13, with its drama and heroic efforts (the flight during which astronaut James Lovell made his now famous announcement) pointed out some important lessons that are appropriate today:

- Human beings are imperfect
- No matter how hard people try to get something done, unexpected events may occur that change the outcome
- Plan for the unexpected

Financial planners are aware that the unexpected can and will happen. Even so, until recently, uncertainty could not be quantified and dealt with objectively. Even worse, in many planning models, uncertainty was ignored. For example, one of the authors recently worked up a retirement analysis for a client using a planning program with both deterministic and stochastic capabilities. In the initial run, the final year of the deterministic model projection showed the client with over $5 million. On the other hand, the stochastic version showed a probability of financial security (having money in the final year of expected life) at only 70 percent. Houston, clearly we have a problem!

For the client analysis mentioned earlier, the answer to the problem of a lower-than-acceptable level of probable success was not apparent at first. To increase the probability of financial success, we took a counterintuitive action: We changed the mix of assets to reduce returns and portfolio volatility, while at the same time slightly reducing spending. The result? A lower final amount of assets, according to the deterministic model. The stochastic model showed the probability of financial security increasing to 83 percent!

How Do I Know I'm There if I Don't Know Where I'm Going?

Critics of stochastic modeling sometimes argue that inaccuracies stem from incorrect assumptions about inputs such as the distributions and averages of future returns, correlations between different investment behaviors, and variances in the client's financial behavior. Their argument is that the uncertainties of inputs are so great, the outcome produces an inaccurate or misleading answer.

We contend few planning professionals have studied the factors affecting a model's accuracy in sufficient detail to make such claims. In fact, most planners are not aware of how parameters that make up the inputs are derived. To really understand model outcomes, the reader must look beyond intuitive logic for answers about accurate inputs. Technically speaking, stochastic model inputs are known as parameter estimation.

How accurate do assumptions and data inputs have to be to obtain realistically accurate answers? Well, if you have ever been in an unfamiliar city or town and have been slightly "lost," you realize the problem stems from the lack of a familiar landmark—a benchmark. In modeling, you must have certain benchmarks, or you won't know if your data inputs are inaccurate—you are essentially lost. In stochastic modeling, the main problem stems from the need to clearly understand the behavior of the investments in the model. This behavior must be accurately established.

To accurately forecast the future behavior of various investments used in a client's portfolio, the planner must have knowledge of the following indicators:

1. The expected return of each individual investment[3]
2. The variance of the returns of each investment
3. Correlation of investment performance among different individual investments in the portfolio and correlation between different years

Let's look at these three indicators separately.

Average Future Expected Return: Living Large

We know there are resources to help financial advisors develop a set of reasonably accurate historical data on almost any publicly traded investment. If these data are "accurate," is everything okay? Most seasoned investment professionals stop at this point and remember the investment-return charts they have shown to clients or potential investors. They all say something like, “Past performance is no guarantee of future performance.” Even if the planner has extremely accurate historical return data for the portfolio, how does that relate to future performance? How can a planner be assured the future will in any way look like the past?

Outside the realm of finance, a large body of empirical research resulted in a concept known popularly as “reversion to the mean.” A well-known and highly analyzed statistic is the long-term return behavior of the domestic stock markets.[4] If there is mean reversion behavior for securities, then periods of excess returns will most probably be balanced by periods of low return. In fact, this is the case.[5] As one small example, growth stocks outperformed value stocks from 1937 through 1968 by 38 percent. From 1968 through 1976, value stocks outperformed growth stocks; 1977–1980, growth stocks outperformed; and 1981–1998, value stocks were once again the reigning champion. For the entire period from 1937 through 1998, the compound returns for growth stocks were 11.7 percent versus. 11.5 percent for value stocks—a tie.

Some still do not believe the mean reverting tendencies of large groups of equities. We suspect the only reason there are dissenters is they either have not studied the research or are unwilling to be convinced for personal reasons. In any case, for large pools of investments such as publicly traded mutual funds, the long-term rate of return of various asset classes has been shown to revert to a mean that is measurable and predictable in a statistically significant manner.[6]

If we stopped right there, we would hear a dissenting voice say, “But that’s not true! The future price of a stock cannot be predicted!” We agree, but the future price of one stock is not what we are talking about, any more than we are talking about trying to predict tomorrow’s average temperature from the previous day’s average temperature. We might get close, but there is no mathematical way to perfectly predict the next day’s temperature. What we are talking about is the longer-term behavior of a large group of securities.

Most important, and central to parameter estimation for a stochastic model, future investment behavior in these large pools can be more accurately predicted if the pool maintains the same underlying investment characteristics. The critical issue is the ability to keep the underlying characteristics of the portfolio the same in the future as it was in the past. This is analogous to a refrigerator. With the proper controls and equipment, the future average temperature inside the refrigerator can be very accurately predicted, because the equipment forces the temperature to stay within the bounds set by the thermostat. We can also predict with a high degree of confidence the range of the temperature around the average.

In the same light, accurate predictions of investment behavior will be difficult or impossible for those whose investment policy is to chase returns by picking the popular stock of the moment or “sector of the month.” Using flawed logic, dissenters to the reversion-to-the-mean concept point at this kind of changing portfolio to argue that future behavior cannot be predicted. As this article does not focus on portfolio design, we will assume financial planners can either construct or find managers who can construct portfolios with investments designed and managed to have consistent long-term investment behavior.

Some planners may wonder about the “new economy” or the “new economic paradigm.” These folks may assume there has been an underlying change in the world, such as “technology makes things more productive,” that will cause mean returns to continually rise. If so, this assumes excess returns will continue ad infinitum. This is not the case. If there is a shift in mean returns due to some change in the underlying economics, it must be verifiable and sustainable over the long term.

The long-term annual average return for large-cap stocks in the domestic markets between 1925–1999 is 11.3 percent.[7] As an example, for a significant permanent change of this number to occur from 11.3 percent to, say, 20 percent, the underlying economics contributing to the shift must change. These underlying economic factors might include lower labor, raw material or energy costs; lower distribution expense or cost of capital; or significantly sustained higher productivity.[8]

Current economic data do not support large enough changes in the aforementioned factors to predict or support this much of a long-term shift.9 Thus, we can assume past data of the mean investment return of portfolios with investments engineered to have consistent long-term investment behavior will provide accurate projections of future mean returns, all other things considered. (See sidebar, “A Long-Term Debate.)

Correlation: All for One and One for All

Correlation is a measure of the degree of dependence or co-movement in the returns of different securities (in formulas, the Greek symbol r is used for correlation). The correlation coefficient is a popular measure of the degree of covariance. Correlation coefficients tell the story of how one security or investment behaves relative to another security or investment over time, or how the investment behaves relative to itself over time. If there is non-zero correlation between two different investments, the behavior of one investment might be used to predict the behavior of the other investment. If there is a statistically significant pattern, they are said to have correlated behavior. In addition, there may be correlation of investment behavior between different years, called serial correlation.

Correlated behavior of securities has been recognized for some time. However, it was the landmark work by Markowitz that recognized that just having knowledge of each individual security’s return and variance was not sufficient to describe the return and variance of a portfolio.[13] While the mean or expected return on a portfolio is just the weighted average of the expected returns on the individual securities in a portfolio, the impact of the variance of an individual security on a portfolio’s variance is a more complicated mathematical matter. In fact, when the numbers of securities in a portfolio are large, individual security returns may have less of an impact than the covariances between securities. While relative magnitudes are clearly important, Fama noted that in portfolios of 20 or more stocks, the variance of the portfolio is primarily determined by the pair-wise covariances of individual security returns.[14]

There are only a few financial planning models at this time that provide for correlation coefficient inputs. Why is this? First, the current programs do not have individual investment inputs, and thus have no reason to accept correlation coefficients. Second, the individual variances and covariances of the securities in the portfolio can be expressed as a number that describes the overall variance of the portfolio.[15]

A model that employs correlation coefficients or summarizes all the individual covariances in a singular variance for the entire portfolio assumes past behavior patterns will repeat in the future. The use of correlation coefficients or portfolio variance estimates come with a caveat, however. There are factors including human behavior, world events, market psychology and capital availability that affect the distribution of returns of individual and groups of investments. As a result, the reader is cautioned that, in order for historical correlation coefficients of the individual investments, securities or asset sectors to persist in the future, the portfolio must have investments engineered to have consistent long-term behavior. Then, current correlation coefficients will provide accurate projections of future co-movement between investments in the portfolio.

Distributions: the Good, the Bad and the Normal

As the debate about stochastic modeling has proceeded, we’ve realized a significant part of the discussion regarding data used for future projections focuses on whether the distribution of past investment returns can accurately predict future investment behavior. More specifically, planners are concerned that distributions in past investment returns may not be consistent with future distributions. This question centers on whether random samples of past events such as the annual return from a stock or group of investments will eventually fit a predictable distribution. The answer depends on the behavior of the underlying securities.

In the early 1800s, mathematician and statistician Francis Galton refined the idea of a normal distribution when examining natural occurrences such as how human genetic characteristics change from generation to generation.[16] His seminal research and publication of distributions has remained a mainstay of statistics for nearly 200 years. His work in normal distributions spawned work in statistical mathematics that is used even today in college finance courses to describe the behavior of security returns.

A large body of research originated in the 1960s to determine the nature of stock return distributions. Results from the 1960s, 1970s, and continuing research on this subject has confirmed the *non-normality* of individual stock returns.[17] Short-term (daily and monthly) returns are characterized both by many outliers and bunching around the mean. This behavior is called *leptokurtic*.[18] Longer-term equity returns have often been skewed and are neither normal nor lognormal. This may come as a surprise to some in the financial planning community.

In Fama's landmark work on security returns, he found that daily returns exhibited leptokurtosis. When he examined monthly returns, the distributions still exhibited leptokurtosis but to a lesser degree. As a result, he argued monthly returns were "close enough" to a normal distribution to claim such returns were, indeed, normal. The computing power available to Fama during his 1970s-era study was just a little more advanced than the computing power available to NASA for the Apollo program. Today, a small PC system easily exceeds the total computing power available to a university or government researcher in the 1970s.

As a result, we no longer need to assume, for mathematical and computational simplicity, that investment returns are normal. In fact, we can model such returns using any distribution—or any algorithm capable of computing a probability distribution. With current computing power we can quickly generate distributions that are not normal, that exhibit mean-reverting behavior and that vary with time. Why should we adhere to a normal distribution of investment returns today merely because it was computationally expedient in the past?

This is important. For example, if future investment returns are not normally distributed but the model assumes a normal distribution for sampling, will the outcome be flawed? The authors carefully examined the distributions Fama used, and found they most closely resemble leptokurtic distributions, which have several interesting attributes.[19] As one of their more interesting properties, these types of distributions have no moments that can be calculated. Put more simply, you cannot calculate a standard deviation for this type of distribution. Is this an important fact? We'll see.

I Was Lost But Now I'm Found! Sort Of…

In the 1960s, NASA was able to make it to the moon with relatively primitive computers because they knew precisely where the moon was going to be. What if the moon's parameters were constantly changing? Could they have made it? Maybe, but probably not without some simplifying assumptions to help get the spacecraft "close."

Today, much more powerfully computational equipment affords us the opportunity to factor in the parameters of uncertainty more explicitly. Simplifying assumptions are no longer necessary for computational purposes. As a result, for parameter estimation, if we have solved the questions of accurate returns, variance and correlation behavior, but use outcomes based on a distribution that only approximates equity return behavior, how accurate can the stochastic model outcome be? Is there a return distribution that better suits stochastic investment performance prediction? Does it matter what kind of return distribution is used? From a mathematical and modeling standpoint, can we compare normal and non-normal distributions to see what changes occur in the financial plan outcome? These questions are addressed in a subsequent article.

Endnotes

1. A detailed description of these statistical tests is beyond the scope of this article. For a statistical primer, the authors suggest Samprit Chatterjee's *A Casebook for a First Course in Statistics and Data Analysis* (New York: John Wiley and Sons Inc., 1995).
2. Glenn Kautt and Lynn Hopewell, "Modeling the Future," *Journal of Financial Planning*, Vol. 13, No. 10, pp. 90–100.
3. In the context of this article, an "individual" investment is a portfolio of one or more underlying investments, which as a whole exhibit a specific and distinct investment behavior. For example, a mutual fund may be described as an individual investment where its behavior is specific and distinct from other investments.
4. Roger Ibbotson and Rex A. Sinquefield, *Stocks, Bonds, Bill & Inflation Yearbook*, (Chicago: Ibbotson Associates, 2000).
5. For an interesting investment related discussion of reversion to the mean, see the lecture by John Bogle, former chairman of the Vanguard Group, "Reversion to the Mean: Sir Isaac Newton's Revenge on Wall Street," Distinguished Lecture Series, MIT Lincoln Laboratory, Lexington, Massachusetts, January 29, 1998 (revised). Available from The Vanguard Group.
6. Eugene Fama and Kenneth French, "Size and Book-To-Market Factors in Earnings and Stock Returns," *Journal of Finance*, 1995, Vol. 50, pp. 131–155, and "The Cross Section of Expected Stock Returns," *Journal*

of Finance, 1992, Vol. 47, pp. 427–465.
7. Ibbotson and Sinquefield, ibid.
8. Do not confuse a change in return caused by a change in the underlying cost structure of production/ distribution with a change to a component of Sharpe's Capital Asset Pricing Model (CAPM). CAPM suggests the return of a stock is equal to, and is the result of, the sum of the risk-free rate of return (assumed to be the T-bill rate) plus a return based on the future growth of the underlying assets, taking into account the implicit uncertainty of that growth (the so-called market risk premium). Using Professor Sharpe's elegant model to explain long-term systemic changes in return may be a theoretically interesting discussion, but it is not central to this article. Systemic long-term changes will not happen because of a change in perception of the inherent market risk relative to other assets. Actually, the P/E ratio stays the same because the market's perception of future risk does not change. For returns to increase, security prices must increase. But if the P/E ratio remains constant, earnings would have to be driven up by changes in one or more underlying economic factors.
9. Robert Shiller, *Irrational Exuberance* (New Haven: Yale University Press, 2000).
10. Shiller, ibid.
11. Ralph Acampora, *The Fourth Mega-Market Now Through 2011* (New York: Hyperion, 2000); and Harry Dent, *The Roaring 2000s: Building the Wealth and Lifestyle You Desire in the Greatest Boom in History* (New York: Simon & Schuster, 1998).
12. Fama and French, ibid.
13. Harry M. Markowitz, "Portfolio Selection," *Journal of Finance*, Vol. VII, No. 1, March 1952, pp. 77–91.
14. Eugene F. Fama, *Foundations of Finance: Portfolio Decisions and Securities Prices* (New York: Basic Books, 1976), Chapter 2.
15. While this article cannot delve into the mathematical relationship between covariances of individual securities and the overall portfolio, Fama's *Foundations of Finance* textbook describes the relationship quite well.
16. Galton's work is described in Peter L. Bernstein, *Against the Gods: The Remarkable Story of Risk* (New York: John Wiley and Sons Inc., 1996).
17. Fama, ibid.
18. Technically, a leptokurtic distribution is one whose coefficient of kurtosis exceeds the value 3.0. Practically speaking, this means the distribution has both a higher central peak and larger outlying "tails" than does a normal distribution. See the following footnote for more discussion on distributions.
19. Interesting leptokurtic distributions include the Student-T and Cauchy distributions. We know Cauchy distributions have no moments and, thus, no calculable standard deviation. As another observation, one of the characteristics of chaotic systems is that they have no calculable standard deviation.

A Long-Term Debate

Currently, there is debate in both economic and investment circles centered on a reasonable forecast of the long-term outlook for domestic and international equity markets. Distinguished academicians, successful private economists and investment analysts differ widely in their appraisals. Some, including Federal Reserve Board chairman Alan Greenspan and economist Robert Schiller, believe "irrational exuberance" led to the speculative bubble that contributed to the implosion of many stock prices during 2000.[10] Others such as Ralph Acampora and Harry Dent claim the equity markets have plenty of room to grow due to demographics, sector rotation and price corrections in a small number of high-flying stocks.[11] Conversely, research by Eugene Fama suggests future returns may be unpredictable in the short term but will revert to a long-term historical mean that is predictable.[12]

Regardless of an individual's personal position on the long-term forecast of equity prices, the application of modern portfolio theory to develop investment recommendations and design portfolios can lessen the threat of substantial downturns, adverse sector/asset class performance, price implosions and such, while buoying up overall returns. Understanding the volatility and average historical returns of assets and their correlation to other investments can help financial planners who recommend investments do a better job.

In short, planners who argue that rates of return, volatility and asset class correlation inputs are not "accurate" or useful in predicting the future are ignoring decades of empirical evidence. Worse still, planners who ignore or deny the existence of volatility may make decisions for clients with horrifying consequences. Planners who select investments need to carefully examine how they develop investment policy for their clients