Monte Carlo Simulators: Are They Worth the Gamble?

BY JIM OTAR, CFP®, CMT

Since the last bear market, more and more advisors are switching from standard retirement calculators to Monte Carlo (MC) simulators to forecast portfolio asset values. What makes the MC different from a standard retirement calculator is that it adds random fluctuations to a steady growth of the portfolio. The user selects a baseline (assumed base growth rate) and a deviation from that baseline. The model then runs thousands (or millions, if you choose so) of projections by randomly varying this deviation. Finally, it reports the range and probability of these projections.

While the MC model is a step forward from the standard retirement calculator, we should not ignore its mechanics. Here are three key constraints and why you should be aware of them.

KEY CONSTRAINT #1. The first constraint of the MC is how it generates randomness. The randomness is generated using a distribution curve. There are many types of distribution curves, such as normal, lognormal, triangular, uniform, binomial, exponential, and geometric, to name a few.

The uniform distribution curve offers random deviations with equal frequency. For example if the baseline is 8 percent and the range is between –16 percent and +16 percent, then the probability of a 15-percent growth-rate projection is the same as a 5-percent growth-rate projection (see figure 1).

FIGURE 1 Typical Distribution Curves

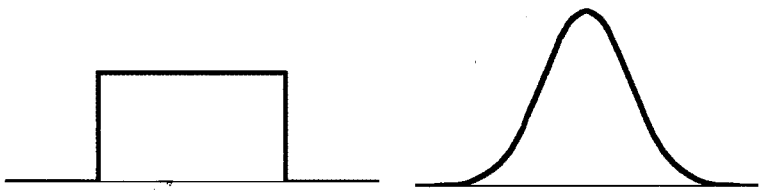

FIGURE 2 Actual Probability Distribution Curve for a Distribution Portfolio

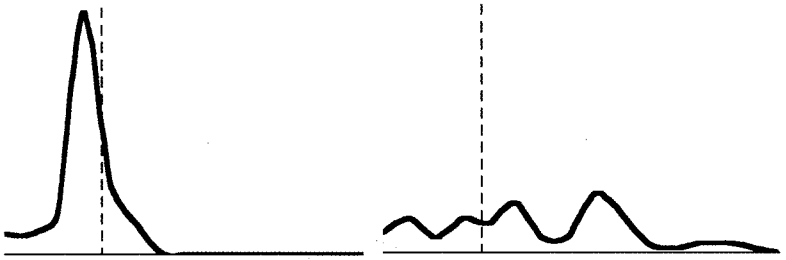

The normal (also known as Gaussian or bell-curve) distribution is based on generating more deviations that are closer to the baseline and fewer that are further from the baseline. For example, if the baseline is 8 percent and the standard deviation is between –16 percent and +16 percent, then a 10-percent growth-rate projection is forecast more often than a 3-percent growth-rate projection.

Historically: An historical distribution curve is significantly different from these idealized distribution curves. Not only that, market history shows that the distribution curve changes shape over time.

Many factors affect the shape of the distribution curve, such as different withdrawal rates, time passed since the beginning of retirement, asset allocation, and rebalancing models. Figure 2 shows the actual distribution curve of a portfolio at five years and 20 years into retirement. As time passes, the distribution curve flattens significantly.

When the distribution curve used in the MC model does not cover the entire retirement time period correctly, the resulting simulations will

be significantly different from actual market history.

KEY CONSTRAINT #2. The second key constraint of MC is that it offers random outcomes and ignores the effects of cyclical (or short-term) and secular (or long-term) trends.

Historically: When we look at history, we observe that markets are random in the short term, cyclical in the mid-term, and trending (up, down, or sideways) in the long term, as depicted in figure 3. In addition, the sequence of market events are not random, they are correlated: Higher inflation eventually causes short-term interest rates to rise, which can have bearish effects on stocks and bonds, and vice versa. Picking growth rates at random for different asset classes and attaching these to a randomly selected inflation rate is not congruent with what happens in real life.

Consequently, the user ends up increasing the range of outcomes, say from ±15 percent to ±30 percent, to cover this constraint of MC. Increasing the range of outcomes only masks the problem, it does not solve it. MC simulation is based on statistical randomness around a predefined straight line. Increasing the envelope of these outcomes does not make it more accurate. If the model does not fit well, then running 10-million simulations instead of 10 does not make it more accurate.

KEY CONSTRAINT #3. The third key constraint of MC is the unrealistic sequence of outcomes.

Historically: Usually during the last one-third of a secular bull trend, good news begets more good news. The index moves higher just because many bet that it will continue moving higher. On the other hand, when a serious bear market starts, bad news begets more bad news, sending the index lower. These situations create "fat tails" on the Gaussian distribution curve.

Markets digest the speculative energy of bull runs in one of two ways: Either a sharp sell-off, like a high waterfall, that may last three to four years (e.g., 1929, the NASDAQ after 1999) or multicycle sideways trends, like a meandering river, that may last as long as 20 years (e.g., 1900, 1936, 1964).

FIGURE 3 Secular and Cyclical Trends in the Market

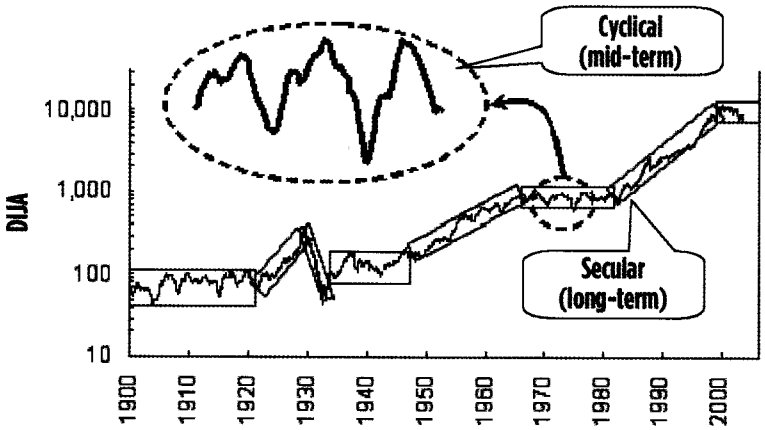

Most MC simulators ignore these as "extreme" or "won't-happen-again" events. They rarely will produce a multiyear, back-to-back streak of multiple bear or bull outcomes, as happens in real life.

Some of the more sophisticated MC simulators do incorporate fat tails. Designers of such models claim that their models can handle fat tails. Unfortunately, they don't. That is because when we look at market history (DJIA or S&P 500), we don't see just one main Gaussian probability curve with fat tails, we see (at least) two, as shown in figure 4.

The leftmost curve represents the "unlikely" sharp sell-offs such as the 1929 crash. During the past century, markets spent about 4 percent of their time in such trends. The tallest distribution curve represents the secular sideways trends, which come with insignificant growth rates and about 5-percent inflation. Markets spend about 50 percent of their time there. Next to it, the secular bullish distribution curve represents an average 15-percent growth rate and 2-percent inflation. Markets spend about 38 percent of their time in secular bullish trends. The right-most curve represents the runaway bullish binges where markets spent about 8 percent of their time during the past century. Keep in mind: I am only reporting what happened during the past century and I make no claims that this century will be the same.

FIGURE 4 U.S. Major Indexes, Distribution Curve for the 20th Century

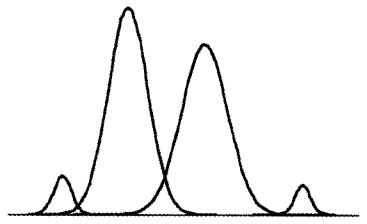

Even if a conventional MC simulator could be designed to incorporate the distribution curves depicted in figure 4, it still would not be good enough. It would be able to simulate the frequency of these events correctly, but it still would miss the sequence of returns. The results from such a simulation would be flawed because, in reality, once markets decide to be bullish, they stay under the bullish distribution curve for as long as 20 years. Then some invisible hand or a seemingly unimportant event pushes the trend either into the left fat tail where it can stay a number of years, or under the secular sideways distribution curve,

>> "MONTE CARLO SIMULATORS" CONTINUED
where it may stay for up to 20 more years. It stays there until the invisible hand or seemingly unimportant event pushes it back under the secular bullish distribution curve.

A well-designed MC simulator should be able to simulate the frequency of growth rates and inflation, as well as the sequence of returns. Whether or not such a model can have any predictive power is, of course, another philosophical question.

The easiest way to demonstrate this key constraint is to look at retirement portfolios through actual market history (figure 5). The median line (where half of the portfolios do better and half do worse) is a lot closer to the unlucky line (bottom decile) than the lucky line (top decile), even though both unlucky and lucky portfolios have the same probability of occurrence (10 percent). In other words, after retirement the path to an unlucky outcome is a lot shorter than the path to a lucky outcome, an effect that no advisor should ignore. Yet we are dishing out retirement plans to our clients using simulators that totally ignore this.

I compared the outcomes of one MC simulator with actual market history using the same case. Table 1 compares the probabilities of depletion.

This only is the tip of the iceberg. What is even scarier is that these flawed MC models are used widely in academic research as the foundation for asset allocation, portfolio optimization, diversification, and risk management.

In the final analysis, most Monte Carlo simulations create solutions that are too optimistic. However, if you insist on using an MC simulator instead of actual historical data, then I suggest that you at least consider using a better MC model, such as the one described below.

A Better Monte Carlo Model

For a better model, you need to start with the bigger picture. Here is what U.S. markets did during the past century (see table 2).

TABLE 1 Comparing Monte Carlo Outcomes with Historical Outcomes

YEARS AFTER RETIREMENT	MONTE CARLO PROBABILITIES OF DEPLETION	ACTUAL PROBABILITIES OF DEPLETION
10	0%	0%
15	1%	3%
20	14%	36%
25	37%	68%
30	55%	86%

TABLE 2 Secular Trends, DJIA 1900–1999

	TREND	AVERAGE ANNUAL DJIA GROWTH	AVERAGE ANNUAL INFLATION	LENGTH, YEARS
All Trends	1900–1999	**7.7%**	**3.3%**	99
Secular Sideways:		**2.4%**	**5.6%**	
	1900–1920	4.2%	4.8%	21
	1937–1948	1.4%	4.8%	12
	1966–1981	0.8%	7.1%	16
Secular Bull:		**15.0%**	**1.8%**	
	1921–1928	20.6%	–1.5%	8
	1949–1965	11.5%	1.7%	17
	1982–1999	15.9%	3.3%	18
Secular Bear:		**–31.7%**	**–6.4%**	
	1929–1932	–31.7%	–6.4%	4
Other:				
Cyclical Bull	1933–1936	33.5%	1.7%	4

Note: The averages indicated in bold are weighted averages by the length of each trend.

A Two-Layer Monte Carlo Simulator

To include the effects of the secular trends, the MC simulator ought to have two layers—let's call this MC2.

The first layer selects a particular secular trend at random. The only rule at this first level is that the same secular trend cannot be repeated: A secular bullish trend can be followed by only a secular sideways trend or a secular bear trend. A secular bear trend can be followed by only a secular sideways trend or a secular bullish trend. A secular sideways trend can be followed by only a secular bullish trend.

The second layer is identical to models in use today. However, they use different base rates for each type of secular trend. I use the following base rates:

If the first layer of simulation is a secular bullish trend, then the second layer is set to a growth of 15 percent annually with a range of ±15 percent, inflation of 2 percent annually with a range of ±1 percent, and a length of time of 20 years.

If the first layer of simulation is a secular sideways trend, then the second layer is set to a growth of 2 percent annually with a range of ±20 percent, inflation of 5 percent with a range of ±2 percent, and a length of time of 20 years.

If the first layer of simulation is a secular bear trend, then the second layer is set to a growth of –20 percent

with a range of ±15 percent, inflation of –5 percent annually with a range of ±2 percent, and length of time of 4 years.

Both the trend type and the stage of the trend are randomly selected at the simulation starting points. For example, the simulation may start in the 6th year of a bull trend, or it may start on the 9th year of a secular sideways trend.

The two-layer simulation minimizes—even eliminates—all three key constraints described earlier. Keep in mind that these particular parameters apply only to DJIA and S&P 500 indexes since 1900. Other markets have different rules based on their own historical experience. You also can change any of these parameters to fit your needs.

The Evolution of the Retirement Calculator Models

Figures 6, 7, 8, and 9 compare retirement projections using different models based on a starting capital of $1 million and an annual withdrawal of $60,000 starting at age 65. All figures are based on holding a balanced portfolio.

Figure 6 is based on a standard retirement calculator with a steady growth rate and inflation. This is the most popular model used by financial planners and it also is available at many Web sites and mutual funds companies. This example shows the projection using an assumed average growth rate of 8 percent and an assumed inflation of 3 percent. It generally is useless for retirement planning.

Figure 7 is a typical Monte Carlo simulation. The probability of depletion in this particular run of simulations was 53 percent by age 95.

Figure 8 is a two-layer Monte Carlo simulation as described above. The probability of depletion in this particular run of simulations was 77 percent by age 95.

Figure 9 indicates outcomes based on actual market history. The probability of depletion in this run by age 95 was 74 percent. The two-layer simulation reflected the historic experience significantly more realistically than the standard Monte Carlo simulator. Keep in mind that these results will fluctuate each time you run the simulation.

FIGURE 5 Actual Market History: Median, Lucky, and Unlucky Outcomes During Retirement

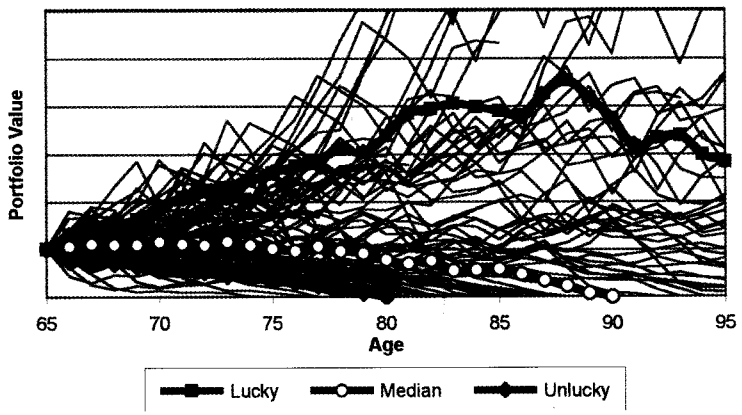

FIGURE 6 Standard Retirement Calculator

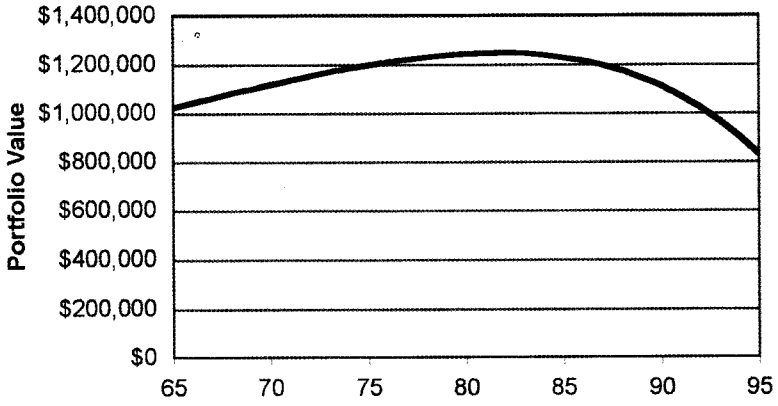

FIGURE 7 Typical Monte Carlo Simulator

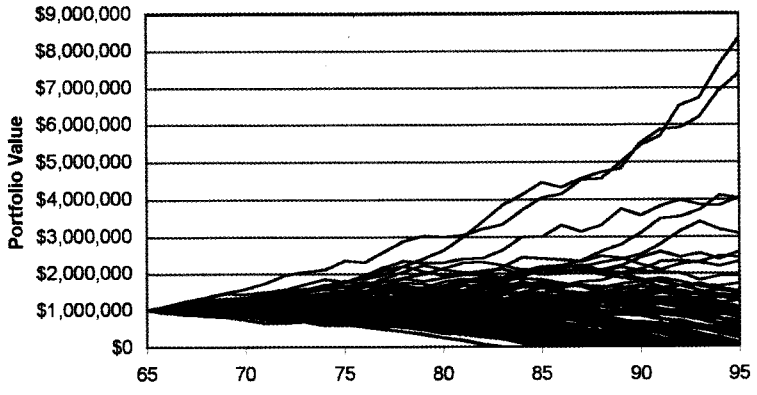

>> "MONTE CARLO SIMULATORS" CONTINUED

Conclusion

Existing MC simulators generally have flaws, which probably are rooted in our own flaws as human beings. Human nature likes observations to fit into a neat, easily explainable, Gaussian mindset. We must go beyond that. Many in the financial industry already know that market events do not fit neat models; rather, they are complex and nonlinear. We must move beyond "projecting 30 years into the future" based on our limited assumptions and simulations. We must give our clients outcomes including the "unknown unknowns." Using better models, or using actual market history, will go a long way toward achieving this objective.

Jim Otar, CFP®, CMT, is a financial planner, professional engineer, market technician, financial writer, and founder of retirementoptimizer.com. He is the author of High Expectation and False Dreams—One Hundred Years of Stock Market History Applied to Retirement Planning. *He earned a B.A.Sc. and an M.Eng. from the University of Toronto. Contact him at jim@retirementoptimizer.com.*

FIGURE 8 Two-Layer Monte Carlo Simulator (MC2)

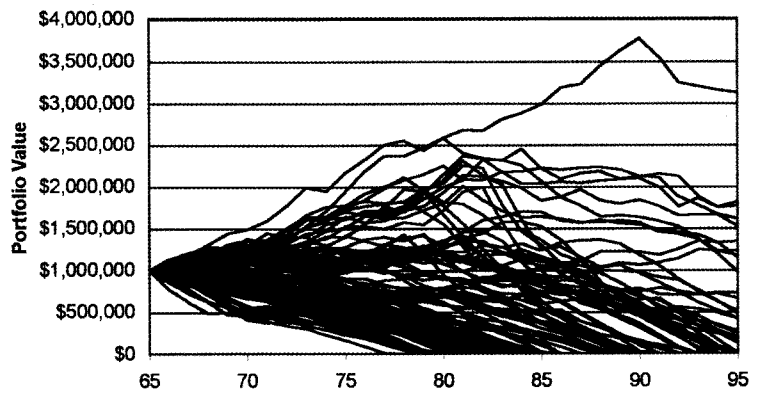

FIGURE 9 Actual Market History