September, 2004

Retirement Income Planning, Part 4*: Beyond Monte Carlo

For the past few years, Monte Carlo analysis has been gaining adherents as the best way to evaluate retirement risk. The Monte Carlo technique is indeed a big improvement over the older, simpler "deterministic" approach. But Monte Carlo and similar models** are being used to answer the wrong question, and so the results they produce are of little or no use.

The essence of the Monte Carlo approach is to create a large number – usually a few thousand – randomly generated scenarios, compute the financial characteristics of each scenario, and then compile the results to determine the overall likelihood of a satisfactory outcome. This is a useful exercise, and helpful to a point, but in the context of financial planning for retirees, it is off the mark. There are four main reasons for this:

1. **The success rates that result are not what they purport to be.**
2. **The results are not meaningful to most people.**
3. **The method fails to address the questions that retirees are actually asking.**
4. **Conducting a Monte Carlo analysis is impractical in key situations**

In this paper, we'll examine each of these four issues, then we'll propose an alternative technique that resolves them. But first, we want to emphasize that we are by no means opposed to Monte Carlo analysis in principle. Our first experience with it goes back to 1981 and was remarkably successful. In more recent years, we have used it to build non-deterministic models of

Monte Carlo analysis is a powerful tool, but like any tool, works best only when applied to the right situations.

---

* Part 1 of this series discussed the urgent and wide-ranging planning needs of people facing retirement, and concluded that if financial companies and employers want to serve this demographic group, they need to address *all* these planning needs. In Part 2 we further explored the follow-up question: can a comprehensive financial planning approach really work for retirees and, if so, how? Part 3 examined investment risks and strategies, and argued that most retirees should be investing conservatively rather than for asset growth.

** Not all "stochastic" models are Monte Carlo models. Some models use modified methods and therefore, strictly speaking, are not Monte Carlo models. A stochastic model recently announced by the Society of Actuaries (SOA) claims not to be a Monte Carlo model at all. However, all of the stochastic models we have seen, including the new SOA model, suffer from the main problems discussed in this paper.

single-point decisions (such as whether to roll IRA money into a Roth IRA). Monte Carlo analysis is a powerful tool, and we are all for it. But like any tool, it works best only when applied to the right situations.

Problems with Monte Carlo "success rates"

Monte Carlo models are being used to predict the likelihood of retirees not outliving their assets. If a computer model could indeed tell us that, the results would be of interest. But that is not what the computer models actually reveal. What they show is the likelihood of the *model* coming out OK.

This might sound like nitpicking, but there is a big difference between any model, and reality. Even a very sophisticated model can only roughly approximate the complexities and unpredictability of real life. The retirement income models that are currently being built are far too simple to claim any close representation of reality. To highlight only the most egregious problems:

- Most risk factors are ignored. The models claim to be "stochastic" (i.e., using variable assumptions) rather than "deterministic" (i.e., using fixed assumptions), but in reality they are only semi-stochastic. They randomize one or two or a few factors in the analysis, but they make fixed assumptions about a wide array of other factors. In real life, nothing is pre-determined.
- They ignore discontinuous and unpredictable risks, such as the possibility that Social Security benefits will change or be eliminated, the possibility that new financial instruments will be invented, the possibility that there will be a radical change in mortality, or that something completely unexpected will occur. Twenty years beforehand, no one could have predicted the Great Depression, World War II, or 9/11. Yet these events had widespread and long-lasting implications. The models assume that nothing unprecedented occurs. Yet the one thing we know from history is that something unprecedented is bound to occur sooner or later.
- Analysis of financial risk is based on data gathered mainly over the past 75 years, as if this somehow defines the outer limits of future possibilities. The fact is that 75 years is not nearly enough of a statistical base to determine the probabilities associated with future financial performance. Yes, it's the best we have, so we have to use it, but let's not pretend that it actually enables us to calculate probabilities for the next 30 years or so. It could be that the 21st century will more closely resemble the 14th than the 20th, and if it does, what use will the current models have been? Similar logic applies to all other elements of the analysis, whether deterministic or stochastic. We simply don't have sufficient basis for determining the range and likelihood of future events.

Monte Carlo models can project, say, a 90% chance of success within the framework of the model, but what is the probability that the model itself corresponds with reality? Probably about zero.

The upshot of these problems is that Monte Carlo models can project, say, a 90% chance of success within the framework of the model. But what is the probability that the model itself corresponds with reality? Probably about zero. This does not mean that the model is worthless, of course. Models do have some semblance to reality. The real problem is that we don't know how much.

So if an actual model predicts a 90% success rate, a hypothetically perfect model might instead predict a 95% rate, or an 85% rate, or a 65% rate. Unfortunately, we don't have any way to measure the divergence between the actual model and reality.

A Monte Carlo analysis might suggest that a certain retirement strategy seems more prudent than another, and that a given plan does (or does not) have a plausible chance of being successful. This is certainly helpful, and for some purposes is good enough. But this is not what the model purports to show, and it could be argued that using a model that can reasonably be known not to do what it promises to do is itself a risky proposition.

Are the results meaningful?

Average consumers, as we all know, do not (and generally would not wish to) understand how a Monte Carlo model works. You would expect, perhaps, that they could at least understand the concept of a 90% probability of success. As a matter of language and simple arithmetic, this is probably true.

But do they really understand what that means? If a plan fails, how badly does it fail, and what will that mean in terms of their quality of life? And what is the difference between a 90% or 95% or 99% chance of success? The studies we have seen indicate that most people have a weak intuitive grasp of probability. So they would surely have no real sense of whether an 80% or 90% chance of success is adequate, or whether they need to pay what can be quite a high additional price for, say, a 98% or better chance. How do we translate these numbers into terms that mean something to real people?

How do we translate "probability of success" into terms that mean something to people? … What people really are asking is: *What should I do?*

One of the consulting firms has come up with the intriguing idea of translating them to bond ratings, so if you want a AAA-rated retirement plan, you need to go for a 99.5% chance of success. This might be enlightening for the tiny percentage of Americans who trade in individual bonds, yet even then, it is doubtful that many of them have a real understanding of the risks associated with different bond ratings. This is the kind of creative thinking that we need in this business – but it doesn't quite solve the problem.

What retirees really want to know

Some retiree somewhere may have spontaneously asked a financial consultant: "If I do such-and-such, what's the percentage likelihood that I will die before I run out of money?" But that person is the exception.

What real people are saying is: "I'm worried, because I know that my resources and options are already pretty limited, I don't know what's going to happen to me or when, and I've never had to face this kind of decision before. *What should I do?*" Furthermore, what they really mean is: what should I do *now*?

Only in the ivory tower world of retirement income model building do people actually try to come up with plans designed to work for the rest of their lives. Real people are usually more realistic. They want a plan that is geared to work well for them under normal circumstances, and that will also make adequate provision for the contingencies that are of the highest concern to them. And they want to know what steps they need to take now to get such a plan going, understanding that adjustments are inevitable along the way.

A model that says: you have an x% chance of success answers only part of this. When the client asks: "X% sounds pretty good, but what if I have to spend my last five years in a nursing home?" the answer had better not be: "That's one of the 100-minus-x% scenarios where the plan fails" (or, worse yet, "The model can't tell you what happens in a particular case").

"X%" is potentially a useful thing to communicate (if it's truly accurate, and if the client truly understands it), but it is not an answer to client's question: *What should I do?*

The "chance of success" form of output is also problematic for another reason. Ironically, the better a model is, the less likely that it can give consistent advice over time. Here's why: imagine that a retired couple receives a report from such a model showing that, given certain assumptions, they have a 90% chance of success. For the next five years, a miracle happens: reality turns out to be *exactly* as the model assumed: income, expenses, rates of return, inflation, and any other assumption in the model turns out to be correct, to the penny. So in five years they return to the same advisor for an update, the same model is run, and now there is only an 80% chance of success. The advisor says they need to make some big changes. What happened? The problem is that they lived another five years, so that their life expectancy is now a higher age than it used to be. Imagine trying to explain face-to-face to clients that, after following the model's advice and having things work out exactly as assumed, the plan now has a higher risk of failure and needs to be revised.

Imagine trying to explain face-to-face to clients that, after following the model's advice and having things work out exactly as assumed, the plan now has a higher risk of failure and needs to be revised.

We are probably better off, therefore, not trying to give retirees the answer to a question that they didn't ask, and that is going to come back and bite us later. Instead, we should focus on answering the questions that they do ask. How to do this is a subject we will return to momentarily.

Problems running Monte Carlo models

As previously noted, Monte Carlo models typically require that a few thousand different scenarios be run. This can eat up a lot of computer time, and there are three ways in which it gets worse:

1. The more precision you want in the results, the more scenarios you need to run (e.g., if you want to try to distinguish a AAA from a AA rating). Even models that rely on two or three thousand runs can produce errors of 1% or more.
2. The scenarios need to be run for each alternative plan. If you are trying to answer a yes/no question (such as, should the retiree annuitize 25% of her assets or not), then you have to run the scenarios only once, or maybe twice (once with and once without annuitization). But if you want to come up with a unified plan where, say, a dozen retiree questions are being addressed at once, each of them with multiple possible answers, the number of combinations is staggering, and running them all through a few thousand iterations is generally just not going to work.
3. As the models become more complex because they try to do a better job of reflecting reality, it takes longer to run each scenario.

Web-based operation has already been ruled out by the developers of most Monte Carlo models. It would bring a web site to its knees.

This problem is exacerbated by the increasing demand that such models be able to run on a website in a reasonable amount of real time, and without bringing the website to its knees. From what we understand, even for the relatively simple models currently under development, web-based operation has already been ruled out . It's just impractical.

We also need to be concerned that the NASD has been reluctant, so far, to allow the use of Monte Carlo models in the sale of securities products.

A better calculation model…

All models are imperfect, so it's easy to criticize. The real question is: can we come up with something better? We think we can, and here are the principles behind it:

- Focus on what retirees should do *now*. This is what they want to know, so let's tell them. Don't try to tell them how much they should withdraw from their funds for the rest of their lives, just how much they should withdraw *this year*. Then let them update their plan every year, so that as things change, or as things fail to change, they can make gradual accommodations.
- Answer *all* their financial questions. If you can't tell them what they should be doing about their pension and about Social Security, about health insurance and long-term care, about adjusting their family budget or even their living situation, about re-allocating their assets and acquiring appropriate new financial products, about what their taxes are likely to be, and so on and so on, then you don't have enough information to give them advice on anything else. All of these things tie together. The aim, therefore, should be to provide an integrated plan (plus maybe one or two alternative plans) that work as a whole and do not leave any major financial issues unexamined or unexplained. As we mentioned above, though, producing a few plans that work may mean examining hundreds along the way.

- Instead of analyzing a hypothetical strategy, or perhaps a few of them, and providing a semi-fictitious "chance of success" report for each of them, justify the suggested plans by showing people the two things they most want to know about: (a) how will they work out if things go as expected, and (b) how can they be adapted if certain adverse scenarios occur? If you can show people that there is an acceptable contingency plan if they live a good long life, or if they need expensive medical treatments, or if they need long-term care, or if their ne'er-do-well child absconds and leaves the grandchildren with them, then that's all they want to know. This is a lot more useful to them than presenting a plan that claims to work 95% of the time. Furthermore, it saves a lot of computer time, because instead of running a few hundred possible alternative plans through thousands of scenarios, most plans have to be tested only once against the "normal" scenario. Just the top qualifiers from that process go through an additional round of contingency plan testing.
- If you cannot produce a plan that meets all of the contingencies that the retiree family cares about, show them how far the best plans do take them, and what would be needed to go the rest of the way. Maybe they will make additional sacrifices now (working longer, living more frugally) to cover themselves later, or maybe they won't. But at least their decision will be based on a clear understanding of the choices they have and the possible real-life consequences, not on mere numbers and percentages that don't translate into lifelike realities for them.
- Accomplish all this by finding out from people what they are worried about, and build adverse scenarios tailored to that retiree and his or her family. Then there is no need to test thousands of scenarios, but only the ones that actually matter to that retiree. Such an approach still contains deep complexities in its implementation, but it is not the computer-killer that Monte Carlo can be, and it is feasible to implement it on the internet.

Is this pie in the sky?

No, but a real-life version of it is still several months away. Still River plans to have a pilot version by the end of March 2005. A prototype of a simple client report can be viewed on our website, www.StillRiverRetire.com, in the RetirementWorks® II section.

If you are interested in learning more, or if you have any comments on the analysis presented here, we'd love to hear from you.