ROLL MODEL

IN 2003, the Social Security trustees began including, as an appendix to their annual report on the system's financial condition, the results of stochastic projections of the system's long-range finances in an effort to illustrate the uncertainty of the results of the traditional deterministic valuation using the intermediate (best estimate) assumption set.

Any time we talk about mathematical modeling, we're talking about some type of system in which behavior can be abstracted in a way that allows meaningful analysis. Let's take the example of a thrift or 401(k) savings plan that currently contains $1,000. Suppose we're interested in determining what the balance would be five years from now, assuming no additional contributions.

Before we proceed, it's important to note that this example *is not intended to demonstrate how the actuaries in the Social Security Administration do their modeling.* Rather, it's intended to provide the reader with an example of how stochastic modeling could be performed to study a much simpler system.

If we were modeling the balance in a purely deterministic way, we would be given (or we would assume) some annual rate of growth, say 6 percent. We would then take the $1,000 and multiply it by 1.06 five times in order to obtain an answer of $1,338.23.

If we were interested in performing what's called sensitivity analysis, which is basically the methodology being used when the term "scenario testing" is employed, we might want to determine what the five-year ending balance would be if the growth rate were -2 percent or 12 percent (instead of the assumed 6 percent). In the first case, we would calculate an ending balance of $903.92; in the second case, it would be $1,762.34.

Now let's take a look at what would be involved in performing a stochastic simulation of the five-year ending balance. "Stochastic" is simply a synonym for "probabilistic." The dictionary says it comes from the Greek word *stochastikos* or "proceeding by guesswork." Mathematically, it's defined as "a process in which a sequence of values is drawn from a corresponding sequence of jointly distributed random variables."

In stochastic simulations, we choose inputs into the model (in this case, the rate of growth) according to some mechanical rules that reflect the probability of that input actually being observed.

Simulating a Savings Plan Account Balance

BY RICHARD M. RASIEJ

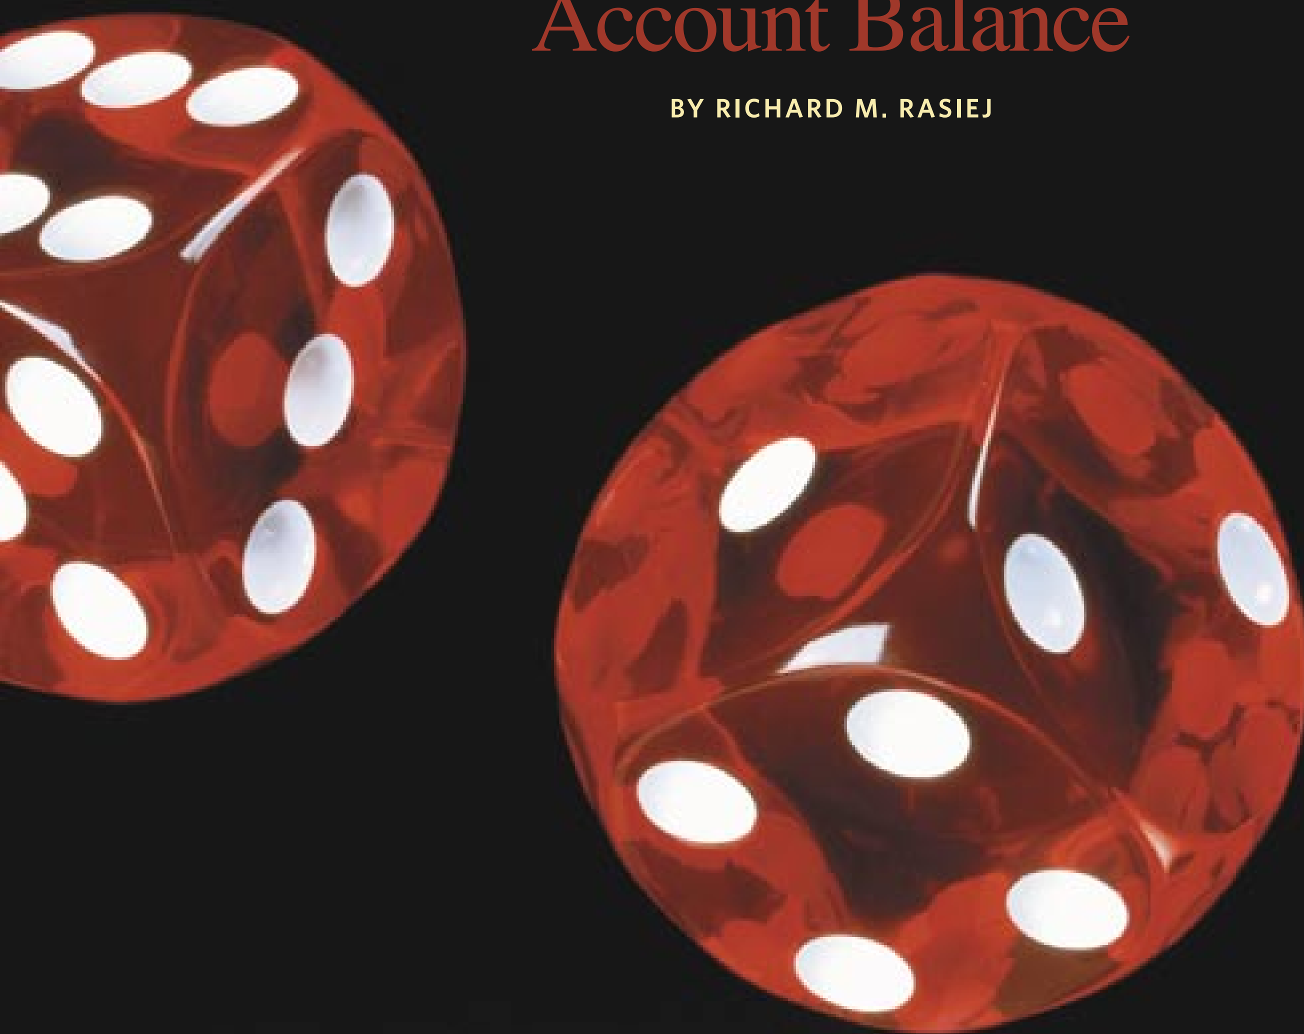

Could stochastic modeling help Social Security actuaries achieve more accurate forecasts than traditional deterministic methods? It's an idea they've been looking at for a few years. Here, in a different context, is basically how stochastic modeling works.

Stochastic Modeling and Social Security

STOCHASTIC (OR PROBABILISTIC) MODELING generally refers to the use of probabilistic and statistical methods to simulate numerous possible evolutionary paths for the important independent variables in the system being studied. This contrasts with deterministic modeling, in which the path of each independent variable is fixed by the assumed values assigned to that variable. The basic tools of stochastic analysis are regression analysis, time series analysis, and mathematical simulation.

The simulation itself (often referred to as a Monte Carlo simulation) is performed by repeating a sequence of trials a large number of times (5,000 or even 10,000 trials are common). Each trial consists of a number of steps:

- Choosing the values to be used for each independent variable (or choosing the amount of random fluctuation to add to an independent variable);
- Running the projection of trust fund balances
- Saving the result.

This process is repeated for a large number of trials. The results of each trial are then tabulated and ordered so that statistical inferences can be made. If, for example, 10,000 trials were run to the year 2030 and the ending trust fund balance was positive in 9,750 of these trials, we would interpret that result to mean that there is a 97.5 percent probability of having a positive balance in the year 2030.

An important potential use for stochastic modeling of the Social Security system is the analysis of various reform options before the U.S. Congress. For example, in the case of recent proposals to partially privatize Social Security through the use of individual investment accounts, stochastic modeling could be used to try to answer the following question: Find an age "X" so that if a worker invests "w" percent of wages into an individual investment account from age "X" to normal retirement age, there is a 95 percent probability that at the worker's normal retirement age, the worker will have combined benefits (from a redefined Social Security reduced benefit and the annuity purchased from the individual investment account) at least as great as under current law.

—Richard Rasiej

Rolling the Die

To restate our problem, let's suppose we wanted to determine the range of ending account balances in a thrift savings plan five years from now, assuming that we start with $1,000 and that we have only one asset class (investment category or mutual fund) in which to invest. We also assume that the only possible annual returns are -4 percent, 0 percent, 4 percent, 8 percent, 12 percent, and 16 percent, each equally likely. We could then simulate each five-year period with five consecutive rolls of a standard die, with faces numbered one to six. We'll interpret a one as indicating a -4 percent return, a two as a 0 percent return, a three as a 4 percent return, and so on.

We can begin the first simulation by rolling the die five times. Suppose the rolls are three, five, one, one, and four. Then, at the end of the first year, the $1,000 has grown by 4 percent, to $1,040. During the second year, the fund grows by 12 percent, to $1,164.80. During the third year the fund grows by -4 percent (meaning it loses 4 percent) to end the year at $1,118.21. During the fourth year the fund loses 4 percent again, to end the year at $1,073.48. Finally, during the fifth year, the fund grows by 8 percent to end the five-year period at $1,159.36.

We would then start the second simulation by rolling the die five more times. This time we roll one, six, two, five, and one. Working through what each roll means produces a five-year ending balance of $1,197.34.

We perform these simulations over and over again (perhaps using a spreadsheet or some other electronic aid), maybe a hundred or a thousand times. After performing all these simulations, we would then have a range of ending values, from $815.37 (rolling five ones) to $2,100.34 (rolling five sixes), with most values in the vicinity of $1,300 (This is because the way possible returns were set up implies an average return of 6 percent per year.)

How would the results be interpreted? Let's suppose that 100 simulations were performed and that the results were distributed as follows:

Ending 5-year balance	Number of simulations
Less than or equal to $1,000.00	12
$1,000.01 to $1,200.00	18
$1,200.01 to $1,400.00	32
$1,400.01 to $1,600.00	15
$1,600.01 to $1,800.00	10
$1,800.01 to $2,000.00	9
Equal to or more than $2,000.01	4

Then, using these results, we could make the following assertions:

- There is a 12 percent chance of ending the five-year period with less than or no more than the starting balance of $1,000 (since in 12 simulations out of the 100 performed, the ending balance was less than or equal to $1,000).

Any time we're talking about mathematical modelling, we're talking about some type of system in which behavior can be abstracted in a way that allows meaningful analysis.

- There is a 4 percent chance of the account doubling or better by the end of five years.
- The likeliest five-year ending balance is between $1,200 and $1,400.

It's important to keep in mind that if another set of 100 simulations were performed, the distribution of five-year ending balances would likely be slightly different. Nonetheless, as more and more simulations were performed, the final results would tend to stabilize.

Often, simulations are performed in models where there's more than one independent variable affecting the system. For example, most people have their savings invested in more than one type of asset. It's usually the case that the various returns on the different asset classes are not entirely independent of each other. Because all asset classes represent available choices in the investment universe, the returns on some classes are influenced by the available returns on others. This phenomenon is called "correlation" is used; it indicates the tendency to move in a similar direction (positive correlation) or the opposite direction (negative correlation).

Raising the Stakes

In order to acquire a feel for how correlation can enter into simulations, let's suppose that we now have two asset classes in which our initial $1,000 account balance is invested, with $500 in each. Suppose that the first asset class is the same as above and that the second is negatively correlated with the first. In other words, if the first asset class does well, the second tends to do poorly, and vice versa. It's possible for both to do well or both to do poorly, although it's less likely.

The following table shows what we'll be assuming for possible annual returns on the second asset class, given a return on the first:

If the return on the first asset class is:	The possible returns on the second will be:
–4%	2%, 4%, 6%, 8%, 10%, 12%
0%	0%, 2%, 4%, 6%, 8%, 10%
4%	–2%, 0%, 2%, 4%, 6%, 8%
8%	–4%, –2%, 0%, 2%, 4%, 6%
12%	–6%, -4%, –2%, 0%, 2%, 4%
16%	–8%, –6%, –4%, –2%, 0%, 2%

In other words, we'll be assuming that if the annual return on the first asset class is 8 percent, then there's a one-sixth chance of a -4 percent return on the second asset class, a one-sixth chance of a -2 percent return, a one-sixth chance of a 0 percent return, and so on.

So now when we perform a simulation, we'll need a pair of dice. The first die (red, to differentiate it from the second, black die) will be interpreted as before. The interpretation of the black die, however, will depend on the number showing on the red die.

For the first simulation, suppose we first roll a four on the red die and a two on the black die to simulate the first year. The four on the red die means that the $500 invested in the first asset class has grown by 8 percent to $540, while the two on the black die indicates that, given an 8 percent return on asset class one, there has been a -2 percent return on the $500 invested in the second asset class, resulting in a drop in value to $490.

For the second year, suppose we roll a one on the red die and a five on the black die. The $540 invested in asset class one will lose 4 percent and drop to $518.40, while asset class two will gain 10 percent and increase to $539. For years three, four, and five, suppose that the pairs of tosses are (three, three), (one, four), and (four, two). Using the information we have about returns of the two asset classes, we'd be able to determine that at the end of the five-year period, the amount in asset class one will have grown to $558.98, while the amount in asset class two will have grown to $581.89, for a total ending account balance of $1,140.87.

As before, we'll perform a second simulation by tossing the dice five more times and interpreting the results accordingly. After performing simulations numerous times, a range of possible outcomes, together with a chart showing the distribution of results (similar to the chart illustrating only one asset class), could be developed. From this chart we'd be able to infer the likelihood of various outcomes.

Here are a few things to keep in mind. First, each of the returns we used in the previous examples does not have to be equally likely. For example, suppose a -4 percent return occurred one-sixth of the time, a 0 percent return occurred one-third of the time, a 4 percent return occurred one-third of the time, and an 8 percent return occurred one-sixth of the time. We'd be able to simulate this situation by using a special die whose six faces were labeled one, two, two, three, three, four and interpreting a one as a -4 percent return, a two as a 0 percent return, a three as a 4 percent return, and a four as an 8 percent return.

Second, all the possible returns don't have to "fit" on a standard six-sided die. If there are "N" possible outcomes, we'll be able to simulate them by envisioning an N-sided die, appropriately marked. (Here is where it begins to be useful to have a spreadsheet or some other electronic aid.)

Furthermore, the number of possible outcomes can even be infinite, as they would be if our annual returns were chosen randomly from an interval ranging from -4 percent to 16 percent. The important thing to remember is that what's needed is a way to sample (simulate) from the "sample space" (the range of possible "rolls" or random numbers) and a way to interpret, in the context of the system being modeled, the value that was chosen. ●

RICHARD M. RASIEJ teaches high school mathematics at Birmingham High School in Los Angeles.