

RECORD, Volume 27, No. 2*

Toronto Spring Meeting
June 20–22, 2001

Session 114PD

How to Create Efficient Simulation Models for Finance and Insurance Applications

Track: Investment

Moderator: CHARLES L. GILBERT

Panelist: PROFESSOR MOSHE AYRE MILEVSKY†

Summary: At this session, participants learn how to create efficient simulation models (in Excel) for various problems arising in finance and insurance.

The panelists discuss:

- *Various applied techniques for simulating long-term equity returns, yield-curve and term structure models and mortality curves, as well as various option payoffs and maturity guarantees*
- *Best practices for the proper stochastic process to use for the variable in question*
- *The appropriate way to 'code' these models in Excel and how to calibrate them*

MR. CHARLES L. GILBERT: Professor Milevsky is going to give a practical and dynamic presentation on how to create efficient simulation models for various problems arising in finance and insurance. Professor Milevsky is the executive director of the Individual Finance and Insurance Decision (IFID) Center in Toronto, Canada. He has a Ph.D. in finance, an MA in mathematics and statistics and a BS in mathematics and physics. The focus of his teaching, research, and consulting work is in the interplay between risk management, personal finance, and insurance.

DR. MOSHE AYRE MILEVSKY: I'm going to be talking about simulations. I'm going to cover a wide variety of topics on a variety of levels. For those of you who do not have a lot of experience with Monte Carlo simulations or have never done them, I'll give you some rationale as to why they are being performed, how they work, and a little bit of the theory behind them. Once we get that out of the way,

*Copyright © 2002, Society of Actuaries

†Professor Milevsky, not a member of the sponsoring organizations, is associate professor of finance at the Schulich School of Business at York University in Toronto, ON Canada.

Note: The chart(s) referred to in the text can be found at the end of the manuscript.

I'm going to give you an example of how to use Monte Carlo simulations to estimate the value of pi, which may seem like it has absolutely nothing to do with actuarial finance or finance and insurance, but in fact the same techniques will be used in some of the other models I'm going to talk about. I will then give you a completely different example of how to use Monte Carlo simulations, examining the relationship between the Euro and the U.S. dollar. I'll give you a rationale for that as well. I'm going to talk about the evolution of cash versus equities from the Monte Carlo simulation perspective, a bit of stochastic calculus background, and finally a case study example number four of the Cox, Ingersoll, Ross model. For those of you who haven't heard what this is, I'll explain it.

What's the idea behind Monte Carlo simulation? Quite often we have problems that are easier to solve by actually randomizing the inputs as opposed to keeping them deterministic. We simulate thousands of different scenarios, and by doing that we obtain an approximate answer. This is very popular right now. In fact, it's incredibly popular among practitioners in the actuarial community, as well as among the public at large. I see individual investors who are now managing their portfolios and thinking, "Okay, what would the Monte Carlo simulation tell me?" It's really reaching the mainstream and is certainly going to be very applicable in the actuarial community.

First let me give you a bit of theory. Imagine we have the following problem. We want to calculate the value of pi, and it's a somewhat difficult problem. People spend many years in the academic community, in the mathematics department, figuring out what the ten-billionth digit of pi is, and this is a very important endeavor, at least from their perspective. What we're going to do is show you how Monte Carlo stimulation can be used to obtain the value of pi. Think of undergraduate mathematics. Forget about the real world for a moment and imagine the following situation:

We've got a square, and a circle is sitting inside of it. We know that the area of this circle is R^2 multiplied by pi. A side of the square is two times the radius. So the area of the square is four times R^2 . What are we going to do with this?

We know that the ratio of the circle to the square—the ratio of the area—is going to be pi divided by four. Clearly, the ratio of the two is going to be the ratio of 3.14159, and so on, divided by four, which is about 78.53%. What does this mean? It means that the circle covers approximately 78% of the square. Why is this important?

Let's say you want to find out the value of pi. What could you do, in theory? It may sound odd at first, but if you want to calculate the exact value of pi, instead of working on analytical techniques to figure out what the last millionth digit is going to be, we can use Monte Carlo simulations. The main idea behind them, and this is the sound bite that summarizes the session, is that the ratio of probabilities approximates the ratio of area. If I know that I'm going to simulate something and I get some sort of ratio, that's going to tell me what the ratio of the areas are, and

areas are associated with integrals, and integrals are associated with expectations, so we can then use probabilities to get expectations. It's ratio of probability to ratio of areas. Then remembering that areas are integrals, and then remembering that integrals are expectations, we've got probabilities to evaluate expectations. What do I mean by this?

Imagine you were sitting here and you started throwing darts at this square and the circle. What percentage of the darts do you expect to sit inside the circle and what percent do you expect to sit outside? It's roughly 78%. We start throwing darts, and assuming we're throwing them randomly, we count the number of darts that are sitting inside the dartboard and the number of darts that are sitting outside the dartboard. That's going to be roughly the approximation for the ratio of the two areas. Think about it. If you throw darts completely at random, many of them will fall inside the circle. A few are going to fall in the edges. The ratio of the two should give you approximately pi over four. Why? Because the ratio of the two areas is pi over four.

For example, if we threw 30 darts and 24 fell inside, then six fell outside the circle. The ratio of 24 to 30 multiplied by four equals 3.2, which is not great, but given that we've only thrown 30 darts, we've got something that's relatively close to 3.14. What have we done here? We're essentially throwing darts and saying that they will land randomly. If they're random, then the ratio of those in one region versus another should be the ratio of the areas. Ratio of probability equals ratio of areas.

If you really want to get the value of pi, this isn't going to help you beyond getting the first digit. You will want to do this many times, and in fact, you can. An equally efficient way to do this is to ask, "Why am I simulating across an entire square? Let me take a quarter of the square which makes it a little bit easier." So that is a big picture—how do I do this in practice? How would I get Excel to generate the value of pi? Obviously, Excel knows the value of pi, but let's say I want to clear the simulation to get the value of pi. What would I do? I know that if I have a square, and each length of the square is one unit, I know that the area is the same ratio of pi over four. We're going to generate two identically distributed random variables. I'm going to randomize a number between zero and one, which is going to represent where on this axis the dart is falling. Then I'm going to randomize a number from zero to one that is going to represent where it falls on the Y-axis. I know that if the sum of the squares of those random variables happens to be less than one, I'm inside; and if it's greater than one, I'm outside; and if it's equal to one, I'm exactly on the edge.

Once again, we're going to simulate two uniformly distributed random variables, each one representing a coordinate for this dart that I'm throwing. Instead of throwing darts, we're going to use a machine, and we are going to have the coordinate fall inside many times, and sometimes fall outside. If, in the simulation, X^2 plus Y^2 is less than one, the number falls inside the circle. If X^2 plus Y^2 is greater than one, the number falls outside the circle.

So here is the formal algebra on how you do the simulation. X is uniformly

distributed, zero-to-one, the X coordinate. Y is uniformly distributed zero-to-one. We are throwing darts at the square and circled dartboard. We define a function. This is exactly what we do in the most sophisticated simulations, as well as the simple ones. We define a function, which is a function of X and Y. If the function is less than one, meaning, if $X^2 + Y^2$ is less than one, we put down the number four inside the circle. If $X^2 + Y^2$ is greater than one, then we put down the number zero. Theory tells us that the expected value of this function is pi. Why? Graphically, this is a circle, and if it happens to be less than one, we're inside; if it's greater than one, we're outside. It's the ratio of the probabilities.

We've simulated 30 random numbers; we have a number for pi, 3.2, which is not very accurate. How do we know if we're getting a good result? We know intuitively that 3.2 is not the value of pi, but in practice how do we know if we got the right result in our simulations? Maybe we made a mistake. How wide are the bands? There are some things you should remember when you're doing a simulation, whether it's to calculate the value of pi or the probability that the yield curve will be inverted in 40 years. The number of trials represents the number of darts that I've thrown. The mean (M) is going to be the mean of these darts that I threw. How many times is it inside? How many times is it outside? What is the ratio? That is going to be the mean. The standard deviation of that series of numbers is going to be S. The standard error of the estimate, if you actually calculate the mean square error of your estimate, take the standard deviation, is S over the square root N.

What does this mean? It means that obviously, the greater N is, the more simulations you throw, the more darts you throw, the greater the denominator is going to be and the smaller your error is, which tells you something that's immediately obvious to everybody. The more simulations I run, the better my number is going to be. It's more than that. It's at a square root N rate, which is an important concept. It's not just that we ran a million simulations so we know we have a better result than running a half a million simulations—that's obvious. It's the amount by which things got better; it's the square root N. One of the things to look at, and hopefully we'll have a chance to discuss, are the tricks you can do so that the denominator is not square root N, but perhaps something even larger than square root N. Why would you want that? You want that so your simulation is more efficient.

Ideally, I'd like the denominator to be N to the power of 20. That's a huge denominator; my standard there is very small. Can we do something with our simulation? Can we do a trick so that the N in our denominator becomes much larger, so we have a better estimate? This is incredibly important, I find, especially when working with practitioners in the actuarial community. There's some magic thing about 500 simulations. When I was in graduate school, I would have failed a course if I did 500 simulations to calculate the value of anything. When I asked people why we are doing 500, they say, "Oh, it's just so time consuming to run these things, it takes us days and days." Well, then do something a little bit more clever with your simulation, so that the denominator is not the square root of 500, which is not a very large number.

More importantly, what's the range for the actual value? We just ran 30 darts and got 3.2 for pi. Where's the real value of pi? That's the estimated value. What's the real value of the growth rate, or what's the real value of that probability? We know from the central limit theorem or actually the law of large numbers as well, that the true value (V) is going to be sitting somewhere between the M that you've estimated minus the confidence interval (S over the square root N), plus the confidence interval (S over the square root N). This means that the greater your value of N, once again the greater the denominator, the more the interval is going to shrink. So if somebody wants a 99% confidence interval for the value of pi, you know how many simulations you have to run, as well as the fact that the confidence interval is going to be mapped via the alpha.

How do we actually get Excel to do this? Interestingly enough, Excel can simulate random numbers. You don't even have to buy a sophisticated package to do it. In Excel we go to the tools menu, data analysis, and random. That allows you to specify how many random numbers you want, any number of them, and what matrix or position you want to put these numbers in. I ask Excel to generate 200 random numbers—100 in the first column and 100 in the second column. Excel wants to know what distribution do you want it from—normal, log normal, or uniform zero-to-one—because we're trying to get the value of pi. It will give you a choice of all the distributions you want. In the first column we have 100 of our random numbers that were generated by Excel uniform zero one—think X coordinate. The second column will be the Y coordinate and it's going to be another 100 numbers. Finally, in the third column, I have an "if" statement, which is as follows: If the sum of this column squared plus that column squared is less than or equal to one, put a four there. If the sum of those two columns squared is greater than one, put a zero there. End of result.

We then add up the third column, which tells us that with an average of 100, we get 3.2. That is the average value of pi, not very good. We know that pi is 3.14. The standard deviation of that series of numbers is 1.6081. That standard deviation divided by the critical square root N, square root of 100, in this case 10, is going to be 0.1608. This means that the true value of pi in this experiment, based on your simulation, is anywhere between 3.2 plus whatever confidence interval you want; say a 95% confidence interval which would be about 3.5 and 2.9. This is not very encouraging.

The message I'm trying to make is 100 simulations gives you nothing in terms of the accuracy of the value. Now, let's do 1,000 simulations. We simulated 1,000 numbers for the X coordinate and 1,000 numbers for the Y coordinate. Then we summed the square and said if it's less than one, we go to four; otherwise, it's zero. What are our results? The value of pi is 3.1. Okay, these two significant digits are encouraging. However, we've got a standard error of .05. With that error, we have values from 3.16 possibly all the way down to 2.9. Is this building any confidence in 1,000? How many simulations do we have to do until we actually capture the real value of pi?

If you have nothing better to do on an afternoon, you can simulate this in Excel. You realize you need about 150,000 of these numbers before you actually get to 3.14. The number of simulations is critical. Otherwise it's just a "what if?" analysis, instead of a simulation.

What if the Dow goes up down, or stays the same? That's not a simulation, that's a scenario analysis. The general application of this, beyond just getting something like the value of pi, is to get any integral that you're interested in evaluating. For example, let's say we have this curve (Chart 1), which I denote with "fx," and that curve takes you from zero all the way to some value, little "t," and for some reason you want to evaluate this integral. This integral might represent the present value of some quantity, or an expectation of some quantity. How would we evaluate this integral? In one dimension it's not going to be a big deal, but in very high dimensions, it becomes somewhat complex. What we're going to do follows the same idea. We're going to throw darts and build a little rectangle outside of the area that we're interested in and we'll say if the darts fall in the area, we'll put a number one there, and if the darts fall outside the area, we're going to put a number zero there. The ratio of those two areas, assuming you know what the area of the rectangle is, which you should be able to calculate, is the actual value for the integral. How many simulations would you have to do? One hundred, 500 and 1,000 are not going to do it. You need to try the tens of thousands before you get an accurate number for this integral. Of course, that depends on how much accuracy you want. Is it basis points or a notional of \$1 billion where basis points aren't going to be enough? The point is integrals—which, as we know from calculus, are areas—can be evaluated by simulation. They are the flow. We evaluate areas by doing a simulation.

Before I talk about the efficiency enhancements when doing a simulation, and what can we do to get that denominator higher, let me ask you the following: If you were trying to get the area under this function, how would you build your rectangle? Do you need a rectangle because you have to know the area? Would you build your rectangle very high up, three or four times the level of the function is? Would you put the rectangle way up there or would you try to put the rectangle as close as possible to the area you're trying to approximate? Intuitively, I'd want to get it as close as possible. If you're going to put this rectangle very high up, it's like having a circle sitting inside an enormous square where you're going to throw all these darts. Very few of them are going to land in the circle. Most of them are going to land inside the rectangle. Your efficiency goes down because you have wasted space.

We'll do the same type of thing with simulations. We're going to try to find the right area in which to embed this curve so you're not wasting unnecessary space. The hint is to get as many tangency points as possible so that we don't waste space and we actually value something that we know is a lot more reliable.

In terms of improving simulation techniques, there are four techniques to get that denominator up. By getting the denominator up, your standard of error is going to be smaller, which means you have more confidence in your estimate. The first method is known as antithetic variables. What are antithetic variables? They exploit

symmetry. You just generated a value for X and for Y. You squared them and you summed them. If the sum is less than one, it's in the circle. What you can do, instead of simulating a whole new pair of Xs and Ys, is reverse them. Instead of having X and Y, call X, Y and Y, X. You're basically taking the dart and symmetrically transposing it in the context of that picture. What's the point of this? When you actually do this, you're saving time because you don't need a million simulations. You only need a half a million. Every time you get those two pairs, you reverse them as well. But more importantly, the correlation between your estimates now becomes a perfectly negative one.

What does that mean? We simulate around a variable and I then invert in order to run another simulation with that same random variable. When I do that, I've got a perfect negative one correlation, which means that the variance of my estimate now has the covariance term of these two simulations, but the covariance term is going to give a perfect negative one on the correlation. My variance and estimate are going to go down, which de facto means that the denominator will go up. That's one efficiency enhancement.

Another efficiency enhancement is the idea of control variance. I like to call that exploiting analytics. The control variance technique means that you try to build the square so it is tangential to your circle. You try to build that square so that you have the least area between what you're trying to estimate and the slack that you need in order to estimate the quantity that you already have a number for. That's essentially what I said earlier; put your rectangle as close as possible around your circle. If you know the value of an integral that's similar to that integral but not quite the same, you know how to integrate E to the X, but you don't know how to integrate E to the X cubed. What you do is say, "I know the value of one of those integrals. I will use that as the control variance and essentially simulate for the E of the X cubed." However, the gap between these two will be much smaller. Again, my variance is going to be reduced.

What is the notion of importance sampling? The idea behind importance sampling is to make sure that when you look at the values that you're generating, you're not wasting the actual value of the simulation on information that you don't need. In other words, if you're actually throwing darts, make sure to focus on the region of the dartboard where the square and circle are embedded. There's no reason to actually go anywhere outside of this area. This is separate from building nice squares as tangential as possible to the circle. When you're throwing your darts, don't waste them in regions that will be of little importance.

With stratified sampling, if you're trying to estimate probabilities of tail events, don't waste time throwing darts at the other side of the distribution. If you're trying to get a ruin probability, if you're trying to get to the 99th percentile, why are you throwing darts in the region of the 70th and the 50th and the 40th percentiles?

I know I just summarized a three-hour lecture in about sixty seconds, but it's important to remember the four terms because these constantly enable you to increase the denominator in your estimate for the confidence interval. I'll give you

some examples of how this works.

There are a few things to remember about Monte Carlo simulation. It's unfortunately become confused with another type of simulation, which is probability-based. If you think in terms of analytics or hierarchy, a Monte Carlo simulation is a sub-set of a probability simulation. Probability simulation means the following: Somebody asks you the odds S&P 500 will fall during the next year. What are the odds that, one year from now, the S&P will be under the current levels? What are the odds that the S&P will finish about 1,500 at the end of the year? You say, "Okay, let me simulate it." Is that a Monte Carlo simulation? You take all the historical data and simulate monthly. You're just looking at a probability simulation to determine. You are not really using a Monte Carlo simulation.

What is the future going to hold? You use real-world data, perhaps you boot strap, you fit the analytic distribution. You're answering a very simple question. What are the odds the S&P will go up to 20%? Did you evaluate any integral where you tried to compute an expectation of some function that you don't have an analytic form? It's a future probability simulation.

Monte Carlo simulation is helpful if you're trying to estimate an expectation or some function, and you're doing the techniques we described earlier. You have a function and you want to know what the area is, or you want to know the value of the integral, and you throw darts in order to compute an area. You take the ratio of the two to get an estimate standard of error. There you've got yourself a Monte Carlo simulation.

If you want to know the odds that the S&P will be above a certain level, the whole notion of confidence interval and your standard of error in the variance reduction techniques don't apply. It's just a matter of getting the right density function to figure out what the future is going to hold.

The two have become confused. So you go to some sort of software package that answers the question: Am I going to be able to retire at age 65? Will I have \$1 million or not? It tells you, here are the odds that you'll have \$1 million when you retire assuming you say this, allocate the asset that way, and retire at this point. Is that a Monte Carlo simulation? They certainly call it that, and it certainly sells. We did a Monte Carlo simulation. It's really just a future probability simulation. Monte Carlo simulation is something much more specific—a sub-set. It involves calculating expectations and functions of expectations.

Another point that I'd like to make about Monte Carlo simulation is that it should only be used when there is no analytic solution available. I know this sounds kind of odd, but if somebody asks you to calculate pi, I don't think you'd use a simulation because you know that for as many significant digits as you need for your purposes, there is an analytic value. Unfortunately, this doesn't happen in the real world when people want to know probabilities of certain events. There is this notion, especially among management, that if you say, "Here's the probability. We're going to lose this amount of money overnight," they'll ask if you did a simulation. You'll say "No,

we have an analytic solution for it," and they'll say, "No. We want a simulation. We want to know the probabilities. We want thousands of scenarios, we want to see all the scenarios." I say, "But we have an analytic solution. Why would I waste time?"

Simulations. You simulate when you don't have an analytic solution, not the other way around. That's a problem. There are many simulations that I find completely unnecessary. You would work analytic values, find the formula somewhere, and then simulate 500 of them. This means that you have an enormous standard of error, and then come back with a result that could have been much better and could have been a point estimate for the actual probability. You're unsure about the input parameters. You're going to have the same problem if you do the simulation as well.

Another point about Monte Carlo simulation, or any simulation for that matter, is that it's impossible to generate a truly random sequence of numbers. There's no such thing as a random sequence of numbers. If you have a random number generator, and if it throws out a number every time an electron is thrown out, you'll get the actual value and calibrate it somehow to a random number, but every other sequence is deterministic. There's a rule taking you from one to another. I remember in the 1970s there was a period of time when the random number generator that some IBM computers had was not actually very random. If you actually plotted the numbers and saw that they were almost on a straight line, it would go from an odd to an even, odd to an even. The folks that designed it thought it was random—odd and even. Behaviorally, that's what you tend to do. If I were to ask you to write down a random sequence, I'm sure you would say, three in a row, "No, I have to go to tails and heads." Those are deterministic sequences, they're not random. They're going to mess up the standard of error. You have to be worried about these things, although nowadays in the generators we don't see that as much of a problem. There is no such thing as truly random.

Finally, remember the difference between quasi-random numbers and pseudo-random numbers. It took me a while to figure out that these are two completely different topics. Quasi-random numbers are an attempt to do what we just discussed—get a better sequence of uniform numbers that will cover more of the plane in higher dimensions. Those are known as low-discrepancy sequences. Pseudo random simply denotes the fact that this is a deterministic sequence. No matter how good the engine running a random number generator, it's always going to be pseudo random. Of course, you get enhancements by the quasi-random.

Let me completely shift gears. Let me give you an application of simulations. I don't know if you follow what goes on in South America, but this has been big business in Argentina. Domingo Cavallo, the finance minister there, has decided that after 10 years of linking the U.S. dollar to the Argentinean peso, they're going to switch. They're going to link it 50% to the U.S. dollar and 50% to the Euro. Why is this important? Obviously, if you're exporting and the U.S. dollar is very high, nobody is going to want to buy your product. Second of all, the link to the U.S. dollar is harming them for other reasons as well. Cavallo came up with a very innovative proposal, and this is where his simulations come in. About two months ago, he said

that as soon as the Euro hits parity with the dollar, they are going to convert and the baskets will link 50% to the U.S. and 50% to the Euro. Right now, the actual value of one Argentina peso is \$1 U.S.

I was asked by a couple of Argentinean fund managers to estimate the probability of that event. What we're going to do is compute the probability that within the next year the Euro will hit parity with the dollar. Now, this is of independent interest. If you're speculating on the direction of the Euro, if you know that you're going to convert over to 50/50, what are the odds it's going to happen? The way Domingo Cavallo has presented it, at least the way I heard him present it, was that he said this is a certainty. Just as soon as it goes to \$1, we're converting. How long are they going to have to wait? Let's do a simulation.

How's the Euro been over the last two-and-half years or so? It has not been very good, for those of you who track this. The Euro debuted at about \$1.17, and it's been on a downward trend ever since. It pierced through the \$1 mark, ran halfway through it, and now it's sitting at close to \$0.90. Over the last half a month it's actually gone a little bit lower. It's sitting close to the \$0.85 level. We used the daily value of the Euro over the last two-and-a-half years or so, and we're going to use this to do a Monte Carlo simulation.

What is the daily change in the value of the Euro? What's the distribution of that? If we take a look over this time period, we have large movements on both ends. What are the statistics of it? Well, the mean daily change is negative five basis points. On any given day the Euro falls five basis points. The actual daily median change is 10 basis points, which is interesting. Half the changes were more than 10 basis points, and half the changes were less than 10 basis points, but the average was negative five. The standard deviation was 16 basis points of change. The maximum change in any given day was 238. There was a jump of 238 basis points and the worst change is -177. The skew was about 0.432.

Why do we need these numbers? We're going to do a simulation and look at this distribution in order to figure out where it's going to go. What's the histogram? (Chart 2) The Excel, or the data tester, tells us that this is not very normal, especially because we have a spike that goes above the peak and we certainly have some fat tails on both the left and the right ends. I'm not very comfortable doing a parametric simulation on this, so I'm going to do it non-parametrically.

What do I mean by non-parametrically? For those of you who haven't done this, I am basically going to have a big hat and in that hat I'm going to put every single daily change that has ever occurred in the Euro. I'm going to put in a little piece of paper that's going to say plus seven basis points, minus twelve basis points, and plus 240 basis points. That very tall one where we have a change of about ten basis points will have a lot of pieces of paper. I'm going to shake this hat and pull out one paper at a time, look at the number, and that will be my daily change for tomorrow. I put it back in the hat, shake it again, and take out a number, and so forth. That's called non-parametric. It's non-parametric because we're not making any assumptions at all about the distribution of the daily change. We're simply saying

that we're putting it all into a hat and drawing them out one by one in whatever order they may appear.

How do we do this? How do we simulate something like that in Excel? We build a simple financial relationship and call it C. Let's say that C2 is going to be whatever the value of the Euro was yesterday plus that random change. C3 will be the value of the Euro yesterday plus the random change. We work through this depending on the number of days, where column C represents today's value that we observe for the Euro, and column B is the overnight change. Once again, column B is that random vector. Column B is the column in which you're going to ask Excel to randomize the numbers in that column. Finally, we use the same tools, data analysis, and random number generator. Of course, you can also use some of the commercial packages to do this.

I have simulated the Euro from April 2001 to April 2002. When I used 260 of the numbers from the hat to represent my business year, the Euro went above \$0.90 for a bit, collapsed to the \$0.80 mark and ended up close to \$0.70 of the U.S. dollar.

Let's try another one. I'm going to do 10,000 of these. Next, one, two, we have another sample path. Did this sample path hit one? No, this one didn't hit one either. Why didn't it hit one? Because the hat that we're taking numbers out of just doesn't contain enough of those positive changes to take me within one year up to the level of one. So we do this exercise over and over again.

On my third simulation, I had one sample path where it was hitting one. How am I plotting these, by the way? I'm looking at the column of the 250 values for the Euro and plotting it in Excel. When you do F9, the F9 will give you random numbers again, so you actually see them on screen.

On the fourth sample path, it pierced through \$0.70 on the U.S. dollar. It was a very bad sample path. The methodology, once again, is to do this 10,000 times free simulation omega. We keep track of the highest, the lowest, and ending values. By omega, I mean the mathematical notation used for one particular sample path. We have to keep track of the highest, lowest, and the ending values in case it hits one and comes back down. It can hit one and if it comes back down, then they've converted and it goes back down. It does not have to be \$1 at the end of the year; it can hit and come back down. The probability of that event is higher, so keep track of the highest. We then count the number of times the highest value is greater than one. That is going to give me a future probability simulation (FPS), not a Monte Carlo estimate. It has nothing to do with Monte Carlo. This is an FPS using the paths to simulate the future.

What is the probability that it's going to hit parity? We took a look at the next 12 months. Here is what the result of the simulation was. The future probability simulation with historical frequencies that I mentioned to you indicates a 6.7% chance that the maximum value of the U.S. dollar to the Euro during the next year will be greater than one (Chart 3). What this means is that from the 10,000 sample

paths that I generated, I counted how many times the maximum point on that sample path pierced \$1. Whether it ended up above or below doesn't matter. I found that it only pierced \$1 about 6.7% of the time. So my analysis, using historical patterns, indicates that there's only about a 6% chance that it will be converted in the next year. Chances are it will not reach parity. The Argentines are going to be linked to the U.S. dollar only. Of course, I prepared this presentation about a month ago. A week ago Domingo Cavallo decided to link it to the Euro even though it didn't hit parity yet, for exporters and importers. For people who are coming in and visiting and working, there's not going to be any convertibility other than to a dollar. In fact, the exporters are going to be subsidizing the importers, but we don't want to get into that.

All you want to know are the probabilities. We had a future probability simulation and we found that most of the time the Euro peaked at \$0.925. Only seven percent of the time did it cross that bluish line, meaning that it ended up piercing the \$1 level along the way. We can do this for a large variety of probability questions.

Now, if we assume a parametric distribution instead of a non-parametric, we simply impose the 60-basis-point standard deviation, and the probability of parity increases 30%. Who cares? This shows how critical your assumption is for the distribution of returns. To use non-parametric, you only have about six percent. If you use parametric and assume symmetry, and get to about 30% higher by about 25 basis points, the parametric versus non-parametric assumption is going to make a very big difference when you're coming up with probability distributions for the future, because the parametric assumption, the normal distribution is going to sample from all possible points, and parametric is only limited to history. It's almost like you're saying history can only repeat itself, and if you never had a return of minus 0.238%, you'll never have it in the future. That's kind of odd. That's the difference between parametric and non-parametric.

This is important because we don't even need a simulation in the parametric case. You can actually calculate the probability that a normal variant will pierce a certain level. The point is, you can actually do it analytically. More to the point, and this is something that you have to contend with, is that if you actually went to the derivatives market when you said that your simulation indicates that there's only a six percent chance the Euro will pierce \$1 U.S., and want to buy an auction that will pay me \$1 as soon as the Euro will pierce \$1, it's a knock-in option.

As soon as it hits the parity, you want to get \$1. You ask them, "How much are you charging for it?" You're expecting to hear a very low number. You're expecting to hear a number that's close to five percent of parity, because you know that it's only about six percent probability and then discounting for the time value of money. There's no way they're going to sell you that for five percent. Why? Your simulations indicate it's a low probability event because they're not using future probability simulations. They're doing Monte Carlo simulations and doing it risk neutrally. They don't care about the real-world drift. They don't care about whether the Euro is on an up tick or on a down tick. They're simulating based on their hedging costs, which are an expectation of what it will cost them to manage this

portfolio. That is a Monte Carlo simulation. If you do a Monte Carlo simulation with their parametric assumption for what it will cost them to hedge, they're going to charge a lot more than that.

Now you know about Argentina, and you know that they have this Euro proposal. But more importantly you see that in Excel, we can use historical data to simulate the probability of future events very easily. You just use that data to boot strap or to do it non-parametrically. You will also see that there's a very big distinction between doing it parametrically and non-parametrically. With each type there is a very high variation of the odds of the event occurring.

I'd like to explain that insurance companies are doing FPSs, and capital markets are doing Monte Carlo simulations. This is why they can't agree on the value of a liability. An actuary says that in 20 years the probability of this event occurring will be close to zero. That's an FPS. Folks on Wall Street wonder how much it will cost to hedge this. That's Monte Carlo. It's going to be a lot more expensive. The further out you go, the greater the variation between the FPS result—what the actuary is doing to calculate reserves or valuation or pricing—and what the individual on Wall Street is doing when he or she says we need to price this in order to hedge it. Wall Street is going to have a much higher number. This is why, when you hedge in the capital market with two capital market instruments, the price is always going to be higher. We can do the same type of simulation. We can simulate thousands of scenarios for the evolutions of stock and cash. The frequency is based on history, but history does not have to repeat itself with just a frequency, and the longer the time arises, the lower the probability of regrets.

I'm going to simulate what the odds would be if you were to invest in the S&P 500 that you'd regret it in 20 years, 15 years, or 10 years. You are saving for retirement and you're not sure if you should put your money in a savings account, fixed annuity, or take a chance and put it in the variable side, at least to the S&P. What are the odds that in 20 years you're going to look back and ask, "Why did I put my 401K in the market? I should have just put it in the bank."

In 20 years, what will make you regret the decision to put your 401(k) heavily in the market? Obviously, if you lose money, you're going to regret it. I would regret it even if I don't make as much as my neighbor who puts it all into safe deposits. It's not that I just want to earn some money at the end of the year; I want to beat the risk-free rate. I love the fund managers who say, "We've never lost money in 20-year cycle." I could put my money under my pillow and I won't lose any for 20 years.

Did you beat the risk-free rate? That's the probability of regret, not earning as much as what you could have received from a risk-free account. Let's do a future probability simulation and see what the odds are that I'm going to regret investing in the S&P. Let's take a look at the risk premium being the spread between what the S&P earned and the risk-free account—the T-bills. What has the risk premium looked like since 1926? I looked at numbers from Ibbotson & Associates. What do we have for a risk premium? Around 1933, we had a risk premium of 60%. In 1975,

we had a risk premium of -30%. The 1929 crash gave us a -40% risk premium. This data is the distribution we're going to use. I'm going to pull numbers from a hat. Let's do the simulation.

What do we get when we actually rank the risk premiums? By running the simulation, I found that a percentage of the time I made less than the risk-free rate. There's a 35% chance that at any given year the risk premium was negative. What does this mean? Let's say you have \$1,000 you want to save and you're thinking about putting it into a bank or the stock market. If you put it in the stock market, there's a 35% chance you will regret it a year from now. There's a 35% chance that the S&P 500 one year from now would be under what you would have received by putting it in the bank. You would get about four percent from the bank. There's a 35% chance you'll make less than four percent. What are the odds you're going to lose money? They are smaller, but that's not the benchmark. What happens if we do this for 10 years? How do we do this simulation? We have a hat with numbers in it, we pull numbers from it, write down returns, and here's what it tells us. If you have a one-year investment horizon, the probability that you will do worse than Treasury bills is about 35%, which is consistent with history. Now, let's take a look at somebody who is investing for 10 years. For 10 years she doesn't plan to touch the money, she's not adding to it, and not withdrawing from it. For 10 years, she won't think about it. For 10 years, the probability of regret, the probability of shortfall is now about 20%. In 10 years there's a 20% chance of that you could be earning less than you would have received from the risk-free rate. There's a benchmark and every year it gets higher and interest rates move. If all you wanted to look at was the chance of losing money, that whole curve shifts down. Once again, I don't think that's the question of interest.

What do we extract from this chart? It would be declining exponentially, meaning that if you're planning to retire in 25 years and you really aren't going to touch the 401(k) money for 25 years, the probability that you're going to end up with less than what you would have earned from putting it in the bank is about 10%. That's still a reasonable risk. This doesn't mean that there's a chance that in 25 years we'll have a market meltdown and you'll get less than what you have now. It earns less than the alternative. How do we get this? We use FPS. Simulate repeatedly and you get what the numbers are going to look like.

Let's talk a little bit about insurance issues. Here's the insurance question that we're going to use simulations for. We've just given a 45-year-old a very interesting guarantee. We've sold him a stand-alone option to annuitize in 20 years at a 180-to-1 factor for monthly income. If you're 45 years old, you might be attaching this to an existing variable annuity, you might be attaching it to a mutual fund, it might be a stand-alone, it might be a rider, but nevertheless, this is what we've just done. We've granted him an option to convert 180-to-1. Of course, if interest rates end up moving against us, or if human longevity increases, or if the adoption rate for this becomes much higher, the lapsation rate is lower and we're going to be in trouble because we've guaranteed a 180-to-1 factor and the factors at that time might be very different.

We're going to do a simulation to figure out what the value of this guarantee would be. What is it really worth? Then you can worry about loading and worrying about setting up how much more or less you charge for it. How much do you charge for this option, and what is the proper reserve for this option in an economic sense? Let's think about the random parts of this. What can go wrong over the next 20 years. The annuity prices and the yield curve can change over the next 20 years. If you guaranteed somebody the ability to convert at 180-to-1 and interest rates plunged and you have a situation like what happened in the UK with Equitable, then you have this ability to convert into an annuity at a very lucrative price, or you might call it a bonus. Interest rates are much lower and you end up with a liability on your hands. That's something that is random. Let's simulate.

What other things are random 20 years down the line? I would say mortality is random. I would say that the Society of Actuaries might decide that we're going to do a new survey of annuitant mortality, and we might end up with a new table that isn't just a scale two projected table or setback. We might end up with a completely new table and find that people are living longer. We might find innovations that enable people to stay on life support to collect their annuity check, in which case they have to worry about it. Mortality uncertainty, as well as interest rate uncertainty, can be factors. I would simulate both mortality and interest rate uncertainty. It's a stochastic hazard rate.

What else is random over the next 20 years? Investment returns are random over the next 20 years. They're random if you attach them to a variable annuity, this guaranteed minimum income benefit, or guaranteed annuitization rate, which is more precise. If interest rates move down, that's a problem; but more importantly, what happens if the market actually goes up very strongly? You now have the obligation to essentially convert this at a fixed rate for a huge pot of money. Think about this. Maybe it's obvious to you, but every once in a while I come across people who argue about this.

If the margin goes up, what kind of exposure do you have? You say, "I don't worry about my guaranteed minimum death benefits (GMDBs) because the markets have gone up." If markets have gone up, you now have somebody with three times the pot of money who wants to annuitize. He is going to say, "You gave me a guaranteed annuitization rate 20 years ago. I want it." So your exposure to the equity market is on the upside. You don't just say, "Here's the guaranteed annuitization rate (GAR). Here's the per unit value." It's going to be a function of the fluctuation as well as the correlation between interest rates and the stock market. It's also the fact that if the person has a big pot of money and gets educated about the benefits of annuities (I assure you that that's going to happen over the next 20 years, certainly because of Social Security discussions), people will become more aware of the fact that if I buy one of these, it pays me for the rest of my life. Sure, annuitization rates now are two percent, which I think is going up. That's not a Monte Carlo simulation. That's a forecast.

For lapsation and adoption rates, how many people are going to lapse? Are they going to adopt this? These are questions that are better answered by actuaries than

by me. I'm just going to focus on the interest rates. How would you simulate interest rates over the next 20 years?

Using the annuity simulation is a possible way to simulate interest rates. I will focus on the distribution of immediate annuity monthly factors 20 years from now. I want to know what they are going to cost. I will simulate the yield curve using a model that's called the CIR 1985 model. It stands for Cox, Ingersoll, & Ross, three academics at a variety of U.S. institutions who developed this model about 20 years ago to show how the yield curve evolves over time. Usually it's used for derivative pricing. It's debatable if you should be using it for real-world future probability simulations. You've got to start somewhere. You start here and then we talk about tweaking it for real-world parameters. How does this model work?

First of all you have to understand that this CIR 1985 model assumes the following: whatever the current interest rates are determines the entire shape of the yield curve. So if interest rates in 2017 are three percent, we know the entire shape of the yield curve. That's one factor model. There's only one factor driving the entire yield curve, and all changes in the yield curve are effectively perfectly correlated. Once again, for those of you who are advanced at this, they'll say, "Come on, one factor model, at least there should be volatility, there should be a long rate." Again, we have to start somewhere.

We can always generalize the CIR model. Imagine that all we're going to simulate is interest rates. I'm going to only simulate the evolution of interest rates over the next 20 years. Short-term T-bills. We're going to take the short rate and use a function to map it into an entire yield curve, but we're only simulating the short rate. Everything else follows from that. So if we know the sample path for the short rate in 20 years, we know what the yield curve looks like, because it's a one-factor model. Yet, with a two-factor model, where one factor was a short rate and another factor was the long rate, you'd have to simulate two variables—short rate and long rate—and use a formula to map them into a yield curve. If you want to simulate the entire yield curve, you're going to have to simulate thousands and thousands of points along the yield curve for the next 20 years and worry about the correlations between them. It's an exercise in futility. You've got to make assumptions about factors. Nowadays, the current thinking in the financial economic community is that three factors should be enough to describe the yield curve.

This is a two-generation-old model in which we look at just at one factor. There should be three factors—short, long, and some sort of volatility or curvature. I'm doing a simulation using a one-factor model, and all we're simulating is interest rates. How are we going to simulate the interest rates? I'm going to assume that the interest rate will evolve in a distributional form. It's going to be parametric and I'm going to use a distribution called the non-central chi-square distribution, in which the limit is gamma. You probably have heard of non-central chi-square distribution. Why in the world am I picking the non-central chi-square distribution to fit interest rates? It's more because of the diffusion that generates that.

Let's remember how the yield curve works, or how the term structure of interest

rates works. Usually when we refer to a yield curve, we talk about a coupon bond yield curve. So we say a yield to maturity on an end-year bond is going to be the applied r that solved this equation for constant r across the denominators. What we're looking at and we're going to focus on is the zero-coupon yield curve, and that's an important distinction when you're simulating. It's much easier to simulate a zero-coupon yield curve, and then build your yield curve from that, than it is to start by simulating a coupon yield curve. You can always add the zeros in order to reconstruct the bond, and nowadays you can trade them separately.

The point I'm trying to make is that when you want to simulate interest rates, you have to decide how many factors to use and what yield curve you're simulating. There are three types of curves. You can simulate the forward curve, you can simulate the zero curve, and you can simulate the coupon curve. These curves, in an upward-sloping environment, will always obey this relationship. That's because the coupon curve is an average, and the red curve is an average of the green numbers. An average of an increasing sequence is going to be lower. The forward curve is going to be higher because it captures the curvature of the zero curve. You have to decide what curve you're simulating. I would say let's go with the zero and build the other curves if we need them from this.

In a downward sloping environment, if you're going to simulate the yield curve, the relationship is going to be reversed once again because of the fact that these are averages.

My objective here is to price a life annuity, or specifically, to look at how the annuity factor (the distribution) is going to evolve over time. Then we go back or refer back to our simple annuity pricing formula. What's the annuity factor going to be? It will be the present value of this probability in the numerator divided by the interest rate (r_i). The r_i in the denominator depends on when this is being used. It's not constant in the simulation. We're going to add up a different rate for a different time period. We're not using one valuation rate, we're using the entire curve, with r_i applicable at that time.

Now, how are we going to do our simulation? This is a little bit more complicated than having a big hat and drawing the historical Euro rates or drawing the historical S&P rates. You have to do something a lot more analytical, because we need to connect the short rate to the entire yield curve. We're going to deal with something called a Brownian motion. This is another three hours that should not be summarized in thirty seconds. The Brownian motion simulates a normal variant and it has four very important properties. One is that it's continuous; it starts at zero, so you can think of it as being able to drive without lifting your pen from the paper. It has stationary independent increments, otherwise known as the Markovian property. It doesn't really remember how it got here, but simply continues to evolve according to a Brownian motion. That means that we don't have mean reversion, at least for the Brownian motion. It has the mean of zero, a variance of time and is normally distributed at any point in time.

Equation 1

The Brownian Motion

- 1 Continuous process that starts at zero.
- 2 Stationary and independent increments.
- 3 Mean of zero, variance of time.
- 4 Normal distribution at any point in time.

$$B_t \equiv N(0, \sqrt{t})$$

$$dB_t \approx \pm \sqrt{dt}$$

$$\Delta B_t \approx N(0, \sqrt{\Delta t})$$

SoA, Toronto Meeting 2001

47

We're going to assume that the change in a Brownian motion is normally distributed with a mean of zero and a standard deviation square root of time. That's going to give us our daily, weekly, or monthly changes in the factor that drives interest rates.

Now, you might be familiar with something known as the Geometric Brownian motion, where you have the dB_t term, which is driving the diffusion, being multiplied by your state variable, but the drift component (dt) is going to be growing.

Equation 2

Examples of diffusion processes
that are driven by Brownian Motions

$$dS_t = \mu S_t dt + \sigma S_t dB_t$$

The Geometric Brownian Motion

$$dr_t = \kappa(m - r_t)dt + \sigma\sqrt{r_t}dB_t$$

The Mean-Reverting Square-Root Process

C.I.R. 1985 (*Econometrica*)

The Geometric Brownian motion is what we're going to use for the simulation CIR 1985, which has a mean reverting term.

Equation 3

Interest Rate:
Steady State Distribution

$$r_T \approx \text{Gamma}$$

$$E [r_T] = m$$

$$SD [r_T] = \sigma \sqrt{\frac{m}{2\kappa}}$$

That is an ordinary differential equation, a one-dimensional, ordinary differential equation. What's it saying? It's saying that the force that you feel, the direction in which you move as a function of time, depends on where r_t is relative to m . If m is higher than r_t , as you're looking up at m , m minus r is positive, so you now have a positive force moving you up. On the other hand, if m minus r is negative, you're looking down at m , it's going to come down. This is known as a mean-reverting interest rate. If interest rates tend to get out of line far away from the historical mean end, they're going to come back down. If interest rates get very low, they're going to come back up. The rate at which that's going to occur will be defined by k . The k parameter will tell you the rate at which it's going to be reverted. This is not a deterministic model. We have to add uncertainty. That's going to be the Brownian turn, the random turn that fluctuates.

You're trying to get back down to historical mean, you're trying to get up to the historical mean, but the Brownian is perturbing it. That's the uncertainty in the economy. That's a one-factor model introduced by CIR in 1985. One of the nice properties of it, beside the fact that it's analytically tractable and that it has intuitive properties, is that the distribution of interest rates become non-central chi-squared. That's the rationale for this.

Why are people simulating interest rates using this bizarre distribution? Why is it normal? It all comes back to this assumption, which is analytically tractable.

What will interest rates look like in the limit? How are they going to evolve? Being that it's gamma distributed, we have an expected intricate of m . This means that we're going to simulate interest rates for the next 20 years. Where do you expect interest rates to be in 20 years? That's m . What do you expect the standard deviation to be? That's going to be square root m over $2(k)$ times σ . Does this relationship make sense to you? Think about it. What's this kappa? It was the speed of mean reversion. It's how fast you come back down, how fast you move back up. If you're moving very quickly back to the mean, then the denominator will be very high, which means that your entire standard deviation is going to be lower. This means that you get to specify both where you expect interest rates to go, your uncertainty about that, and how fast it's going to move back to the level. You have three free parameters here, not three factors.

I told you earlier, we're not simulating every single point on the yield curve. We're simulating a factor, which is interest rates, and we're going to use a formula to map that factor into the entire yield curve. The price of a zero-coupon bond $P(T)$ that matures in three years is worth the following: $A(T)$ times E to the power of $-B(T)r$. One of the great benefits of this model, despite the fact that it doesn't fit very well, is it has a tremendous pedagogical insight into term structure models. You want to start here. Let's take a look at that formula PT .

Equation 4

Equilibrium Price of the Zero Coupon Bond (ZCB):

$$P(T) = A(T) \exp\{-B(T)r\}$$

$$A(T) = \left[\frac{2\gamma \exp\{(\kappa + \gamma)T / 2\}}{(\gamma + \kappa)(\exp\{\gamma T\} - 1) + 2\gamma} \right]^{2\kappa m / \sigma^2}$$

$$B(T) = \frac{2(\exp\{\gamma T\} - 1)}{(\gamma + \kappa)(\exp\{\gamma T\} - 1) + 2\gamma}$$

$$\gamma = \sqrt{\kappa^2 + 2\sigma^2}$$

SoA, Toronto Meeting 2001

50

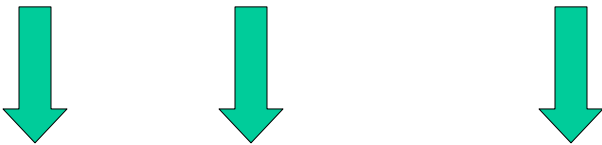
Let's say A is equal to one and B is equal to one. What would we get? P(T) is equal to one E to the -r. That represents the present value of a dollar that I'm going to receive in one year. That's the price of a bond. If B(T) is equal to T, what's the formula going to give me? The present value of a bond issued in two years. So this structure of P(T) corresponds to what you'd expect to see if you had constant interest rates and a flat yield curve. Because of the fact that the yield curve is not flat, and this is a function of time, the A(T) and the B(T) have their own functional forms. The gamma parameter is just the sum of the mean reversion and the volatility. That's not a new parameter. We said this is the variant and this is the speed of mean reversion.

We've simulated an interest rate and built a yield curve from the interest rate. Why is this formula the right one? It has something to do with the equilibrium in the interest rate market. I can't get into it right now, but let's run our simulation. Let's run a simulation in Excel that generates a yield curve using this particular specification. This is how we're going to do it.

Equation 5

Simulating the C.I.R. 1985 Diffusion: Monthly Approximation

r_i	ΔB_{i+1}	Δr_{i+1}
r_0	$N(0, \sqrt{12})$	$k(m-r_0)/12 + \sigma\sqrt{r_0}\Delta B_1$
$r_1 = r_0 + \Delta r_1$	$N(0, \sqrt{12})$	$k(m-r_1)/12 + \sigma\sqrt{r_1}\Delta B_2$
$r_2 = r_1 + \Delta r_2$	$N(0, \sqrt{12})$	$k(m-r_2)/12 + \sigma\sqrt{r_2}\Delta B_3$
$r_3 = r_2 + \Delta r_3$	$N(0, \sqrt{12})$	$k(m-r_3)/12 + \sigma\sqrt{r_3}\Delta B_4$



We've got three columns. The first column is r_i , represents the interest rate at any point in time. The short-term rate that's being used to generate the entire yield curve. The second column is the uncertainty. It's the factor that drives things, it's the Brownian term, which is going to be normally distributed with a mean of zero and a standard deviation of square root 1 over 12. Where did 1 over 12 and the square root come from? We're going to build this monthly. You can build it weekly. You can build it daily, but once you decide, remember that the standard deviation is the square root of that time period. Since we're going monthly, we're going to get 1 over 12. So that's the second column. Once again, that second column is the random one. The third column takes the initial position of interest rates (column one) and the random components (column two) and generates the new interest rate, next period's interest rate. Next period's interest rate is going to be determined in the following way. It's driven or built by the diffusion, by that equation I showed you: K , speed, times m minus r_i . That gives me the direction in which I'm going. It's 1 over 12 because we're only moving one month, plus the uncertainty term, which is going to be multiplied by the Brownian term. This moves us one month at a time for the next 20 years.

Before I move on, what's the takeaway from here? They're all going to be changing on the over time, all of them—every time you hit F9 (the 'calculate' key in Excel). But where are the random numbers being placed? They're in the second column. Everything else is driven from it. If you want to do a three-factor model, to really get sophisticated and say we have uncertainty about interest rates and it's the

curvature as well as the long-term rates, then you'd have three columns in Excel, each with the random numbers generated and you'd have a formula that maps those three columns into the yield curve. It would be much messier, but it's the same idea. Factors are simulated and then mapped into the yield curve.

Finally, let's go over the simulation (Charts 4 and 5). We start with the current yield curve in the short rate. We are going to imply these parameters from the current yield curve, the kappa, M , and sigma. We're going to let the yield curve tell us what those parameters are, or we can estimate them historically, but then we'd have to worry about the implied preference for risk. We're going to simulate this CIR diffusion for the next 20 years. We will have every single one of these paths generating an entire yield curve. Once we've got a yield curve, we generate an annuity factor 20 years out. The final step is to plot the histogram of the annuity factors. Chart 6 is a histogram of annuity factors 20 years from now. Let's compare it to the guarantee that's been provided.

We've simulated annuity factors 20 years from now. These are obviously conditional on survival for 20 years, or taking into account the probability that they may not even make it to 20 years. Of course, we didn't deal with lapsation or surrender. Just what will the annuity factors be for this group of business? What you see is that the number exceeds 180 73% of the time. I didn't pluck that 180 out of thin air. That 180 was actually embedded in a certain prospectus for a guarantee annuitization rate. My point is that there's a very good chance that the annuity rate is going to be in the money. There's a very good chance that 20 years from now you will have to actually settle an annuity factor that's much better for them than what you ended up guaranteeing, which means that you'll either have to give them the difference or the annuity at the old 180 rate. Why is this number so high? Clearly, it doesn't account for the fact that people are going to leave and not take this up. If we have three percent adoption; 97% of the people are going to be subsidizing it. That's not what my discussion here was about. My discussion was how would we simulate annuity factors, and I hope that that's what I've managed to achieve.

I'd like to share some final words. Monte Carlo simulation is an efficient method for computing values when analytic techniques are not available. If you have analytic techniques available, why are you simulating? The future probability simulation is based on the notion that frequencies repeat themselves. That's what we're saying. History repeats itself, but not in the same sequence at which history occurred. We're constantly asked, "Are you saying history is repeating itself?" I'm saying that the frequency with which history repeats itself is constant, not the sample path. We're not drawing an entire sample path from history out of a hat. We have the sample path. We're drawing out monthly or weekly returns. It's just that those are coming from a certain distribution. Results are very sensitive to models used. This is so much so with simulations. You can have one model giving you a 30% chance, and another one a 5% chance. There's a very wide variation. I understand how some of the old timers who look at some of this stuff ask, "Is this really better than deterministic modeling? Is this really better than saying assets are going to grow at this rate or assets are going to grow at that rate? There is such a wide variation. You have a bunch of assumptions going into it and you're not even clear about

when this is going to fit out of sample." I make my career and my job doing the other part, but I have some sympathy for the notion that there is a huge variation. In some sense, unfortunately, it's almost as if you know in advance what kind of results you can get by making sure you're doing modeling assumptions.

We assumed gamma. That's where we got those results. Is that why you got those results, or is that what you wanted to get when you assumed gamma? It worries me a bit and I'm getting more skeptical over time. Running 500 is not enough. Whether you're doing simulations or not, this is a scenario analysis, not a simulation.

What if you want to read more about this? I would say to start off with the following. First of all, Phelim Boyle, University of Waterloo (1977), was really one of the first people who applied Monte Carlo simulations to options pricing. It's actually relatively easy to get through because it's not encumbered by all the sophistication that's come along ever since. It's just trying to explain how you would use simulations.

If you want to learn in more detail how to build these simulations, then there's a book by Bratley, Fox and Schrage, (1987) called, *A Guide To Simulation*. It's a good reference book if this is an area you're moving into. The model that I talked about is in the journal *Econometrica*.

Finally, if you want to learn more about stochastic differential equations, I can give no greater endorsement than the book by Oksendal, which is updated almost yearly. It explains in really simplistic terms how to build a diffusion process, what it means for something to be stochastic, and what it means to have different factors. It's really an excellent book. We use it for both our graduate and undergraduate textbooks in simulation and stochastic simulation.

FROM THE FLOOR: When kappa was equal to zero, why didn't it just converge to mean zero?

MR. MILEVSKY: There's no force in mean reversion to bring it back. If kappa is zero, then even though they're above the mean, there's nothing forcing them to go down, because it's a zero number.

FROM THE FLOOR: How can you generate pseudo-random numbers, and what is the maximum length of cycle you can generate?

MR. MILEVSKY: I'd say that a good number of generators is about 30,000 before you have to start worrying about repetition. The recent ones that combine two generators and try to put the 'if' statement in there to worry about the cycling, are about 20,000. However, if you're going to generate 100,000 of these, you have to worry about them, which is why something like quasi-Monte Carlo is the way to go. You can also buy off-the-shelf packages that worry about that for you. They just generate the random numbers and you use them. After awhile, you get to see the same numbers over and over again.


FROM THE FLOOR: Can you give an example of how you've applied some of your efficiency techniques to reduce standard error in these Excel simulations?

MR. MILEVSKY: I'm trying to estimate an expectation, and that is really something that comes up when we're looking at option pricing. So let's say you want to know what the value of the right is, but not the obligation to get a certain payment only if a certain level is breached. One of the easy answers to your question would be to simulate symmetric random numbers. Every time I've generated a zero one random variable and I've generated a random number from that, I actually take the opposite, the symmetric opposite. That reduces my variance almost by a factor of two automatically. That's one way to go about doing it.

Another way, when I'm looking at the valuation of options, is to find an option that I know about and use that for the control variant for the option I don't know about. For example, when you're trying to price an equity-indexed annuity that pays you an average of some rate or an average of some index, unfortunately we know that the arithmetic average of these numbers isn't analytically known to be log normal. The arithmetic average of these log normal variables is not log normal. The geometric average is log normal. We know what the geometric average would be, meaning, if you were offering an equity-indexed annuity (EIA) that offered a geometric average, it would be very easy to price, and you get an analytic solution. You use that as a controlled variant for pricing the arithmetic option. You've got an efficiency enhancement in these types of cases. It's very much on a case-by-case basis.

Chart 1

General Application




SoA, Toronto Meeting 2001

13

Chart 2

Histogram (Frequency) USD/EUR



SoA, Toronto Meeting 2001

20

Chart 3




Chart 4


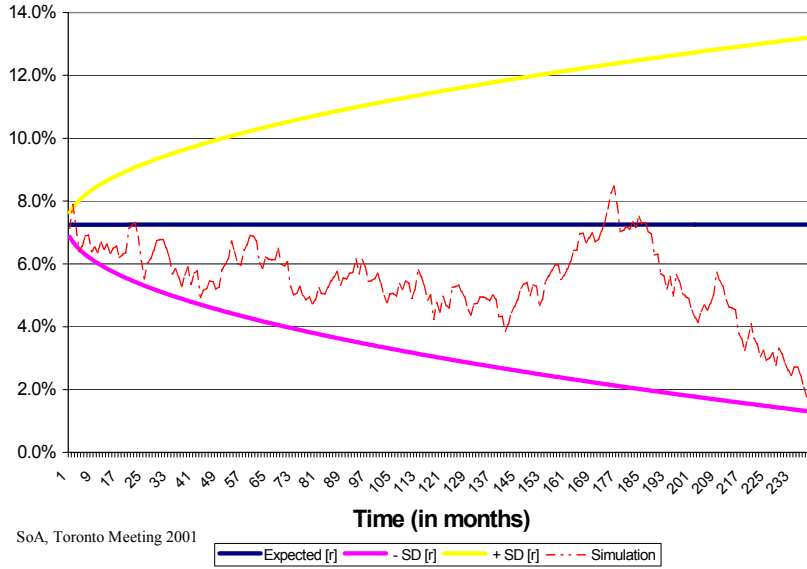


Chart 5

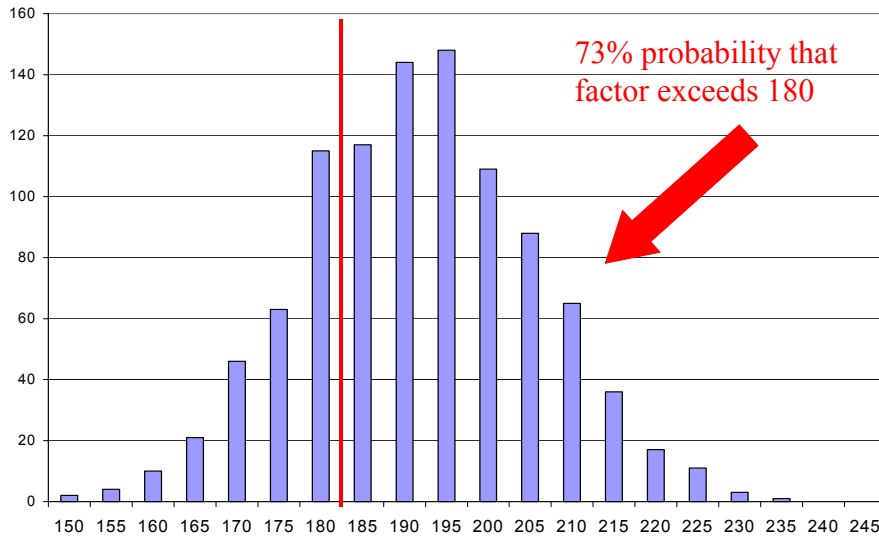
Sample Path and Range for Interest Rate (CIR 85):
 Assuming: $r_0 = 7.25\%$, $m = 7.50\%$, $\kappa = 0$, $\sigma = 5\%$



53

Chart 6

Age 65 Annuity Factors: 20 Years F.P.S.



SoA, Toronto Meeting 2001

55