RECORD, Volume 29, No. 1*

Washington, D.C. Spring Meeting
May 29–30, 2003

Session 68TS
Stochastic Modeling in the Financial Reporting World

Track: Financial Reporting

Moderator: ROBERT W. WILSON
Panelists: RONALD J. HARASYM

Summary: The risks of variable annuities and the coming of international accounting standards have brought stochastic modeling out of the back rooms of investment quantitative analysts and into the everyday world of risk management and financial reporting. Panelists discuss the basics of stochastic approaches to setting reserves and embedded values. Attendees leave with a greater knowledge of how the coming stochastic approaches to valuation will impact financial reporting.

MR. ROBERT W. WILSON: Ronald Harasym is assistant vice-president of Financial Risk Management in the Corporate Risk Office with Sun Life Financial in Toronto. In this capacity, he is responsible for monitoring, quantifying and managing specific capital market risks for the world-wide operations of Sun Life Financial. He is a Fellow of the Canadian Institute of Actuaries and a Fellow of the Society of Actuaries. He is also a Chartered Financial Analyst. Ron is a graduate of the University of Toronto, with an MBA from the Rothman School of Management. He has fourteen years of experience in the insurance industry, and over the past eight years, Ron has worked extensively on the quantification and hedging of embedded option risk in the United Kingdom and the United States. He's a member of the Society of Actuaries' Course Seven Education and Examination Committee and he frequently lectures in the Department of Statistics at the University of Toronto.

MR. RONALD J. HARASYM: I first started working on stochastic modeling back in 1996 where I was a one-person shop looking at embedded asset and liability options. At the time, the embedded options were not being properly priced—they were more or less being given away for free in Canada, the United States and the

U.K. I started applying stochastic models to estimate the cost of embedded options, but not to the liking of the marketing people when I started implying that the embedded options were not without cost to the company. Basically, it started off fairly small, looking at investment performance guarantees on segregated (separate account) funds that were in the pipeline, various book value wrappers, as well as pricing embedded options that were on the asset side.

In 1998, I was asked to look at the risk inherent in guaranteed annuity options on deferred annuities in our U.K. business. This was an industry-wide problem. The risk was similar to the risk inherent in guaranteed minimum income benefit riders that were (and still are) offered on variable annuities in the United States. The difficulty was that no one was looking at the risk of the embedded option and stochastically modeling it. The U.K. was a very deterministic modeling environment. My objective was to model the embedded guaranteed annuity option and to determine an economic risk profile of the guarantee. Ultimately, we used the tool to hedge the risk and evaluate trade-offs between various investment strategies. From this perspective, the use of stochastic modeling worked quite well. My presentation today will describe how stochastic modeling has moved out from the back rooms of offices and into the highly visible financial reporting world.

I want to give you a high-level view of stochastic modeling in a generic framework. It's a framework that I go through when performing any stochastic modeling task. People tend to think of a stochastic model as a single model, but in fact there are often several other processes behind it, such as random number generation and economic scenario generation. For demonstration purposes, I'll use a guaranteed minimum income benefit (GMIB) rider. I'll show the modeling results and some sensitivity testing, comment briefly on reserve and capital relief, and then share some final thoughts.

A few years ago, there was quite a discussion on the discussion boards of the Canadian Institute of Actuaries as to a definition of stochastic modeling. The word "stochastic," courtesy of dictionary.com, is derived from the Greek word *stokhastikos*, which means, in short, "to guess at." You might want to call it something like "intelligent guessing." A stochastic model by definition has one random variable in it and specifically deals with time-variable interaction. The stochastic model itself doesn't have to be a simulation; it just has to have a random variable in it. What we typically do is repeat those random elements, ending up with a series of results. This is what we refer to as a "Monte Carlo" simulation.

A stochastic simulation is an imitation and simplification of a real world system. People tend to think of a stochastic model as the real world, but let's face it, it's just another tool in your tool kit. One of the advantages of stochastic modeling is that if you're going to either offer products that have embedded options or purchase assets that have embedded options, you can try to price the embedded options and assess the risk. Then once you've priced the risk, or at least quantified the risk, you can work the theoretical cost into the pricing. In the past, the majority of these embedded options were never priced appropriately, often being given away for free, which has come back to haunt numerous insurance companies.

Stochastic simulation is a useful tool for forecasting purposes. There are also advantages from a financial reporting perspective. You can forecast where you will be at future points in time. One benefit is that you end up with distributions of results. As long as you have a robust model and a robust framework around it, you can project financial statements. If it's not required yet, it may be required soon; there are certainly advantages to performing stochastic modeling over not doing it.

From a risk management perspective, some people say that if you can't quantify the risk, then it's not a risk. Or they think that what you don't know won't hurt you. But in many products, there are significant embedded options. Unless you quantify these, you don't really know what's coming down the pipeline.

Another use of stochastic modeling is for the simulation of very complex systems where a simple closed-form solution does not exist. For example, if you don't have a simple formula or equation to price your embedded option, then you can perform a stochastic simulation and derive a theoretical or an expected answer. However, keep in mind that stochastic modeling is part art, part science and part judgment. You have to use common sense.

Stochastic simulation is not a magical solution. You need to perform reality checks during the modeling process and understand the limitations of the model. What you get out is only as good as what you put in. If the stochastic model is weak and doesn't capture the appropriate variable interactions, then don't expect the interactions to fall out afterward.

In a stochastic simulation, complex situations with long time frames can be compressed into a more manageable package. You can get a better understanding of the dynamics of your product. You can test policyholder behavior to see where you get hurt the most. People tend to think about the downside risk in stochastic modeling. However, there's also upside risk. The advantage of stochastic modeling is that the whole distribution of risk can to be quantified and examined.

Stochastic simulation is preferred over deterministic modeling when regulations provide real economic incentives, such as significant reserve or capital relief, for performing stochastic simulation. When risk is modeled deterministically or if you model risks independently, then you're not going to pick- -up the benefit of diversification. Modern investment portfolio theory says that for any risks that have a correlation of less than positive one, by combining these risks, there will be some benefit from diversification. When you model risk stochastically and inter-dependently, the end result is a superior risk assessment with the quantification of the benefit from diversification. There are some exceptions when dealing with non-linear risks such as GMIBs. Finally, my favorite benefit from stochastic modeling is that you can watch your company fail over and over again as you simulate those embedded liability options that many people originally deemed to be worthless!

There are limitations to stochastic modeling. It requires an enormous investment in time and expertise. It is technically challenging and computationally demanding. Often, reliance is placed (and companies become dependent) on a few "good" people. Frequently, stochastic models become a "closed shop." It seems as if those individuals who have all the knowledge don't want to share it, don't want to document and don't want to explain their work to others. So you end up with some very challenging situations.

Another limitation of stochastic simulation is that the use of thousands of scenarios may create a false sense of precision. Some people tend to look at a stochastic simulation and automatically think it has both accuracy and precision, whereas in reality, it's only as good as what you've put into it. If your model is specified incorrectly, then your risk assessment will be incorrect and you will be hedging the wrong target. In the end, the stochastic model will become a disadvantage rather than a competitive advantage.

The results of a stochastic simulation can be difficult to interpret, especially to senior management, who might be used to seeing only one number. Often the results can be presented in many different formats, such as scatter plots, histograms, conditional tail expectations and so on. Effective communication of the stochastic results is a challenge. Some people tend to use jargon that is not readily understandable. I find that in communicating stochastic results, you have to keep it very simple and distill the information down to a few basic points.

You wouldn't want to perform stochastic modeling for everything; it wouldn't make sense. There are a number of situations where stochastic modeling is preferred over deterministic modeling, such as when you have skewed or discontinuous distributions or cost functions. Examples of those would be when modeling investment guarantees on segregated funds in Canada, guaranteed minimum death benefits (GMDBs) and GMIBs in the United States, or any situation where there's significant volatility or sensitivity to the initial starting conditions. A good example would be if you have interest rate options, and the current market conditions are materially different than they were in the past.

Where there's path dependence, such as on guaranteed annuity options in the U.K. or GMIBs in the United States, the level of policyholder annuitization will depend upon what path or what level interest rates have been over the period of time being projected. You can capture this relationship in a stochastic model; it may not capture as well in a deterministic framework. Finally, cases where volatility or skewness of the variables is likely to change over time can also be incorporated into a stochastic simulation.

Now I am going to work through the steps that were followed in order to perform the stochastic simulation of a GMIB rider on a variable annuity. A generic framework, from a process perspective, is shown in Chart 1. There really isn't a starting or an ending point. It's a constantly evolving process; you're constantly feeding back and looping through it. In the flow chart, the rectangular boxes indicate processes or models, while the parallelograms indicate input/output (typically data).

Within this framework, model and data validation are very important steps. Some people take the historical data "as is." Yet I've spoken to other people who say they spend an enormous amount of time going through the data, particularly if they're trying to fit fund returns to a benchmark. When measuring fund returns, statutory holidays in Canada, the United States or the U.K. don't always line up, so there may be data lags. Sometimes the fund values are not correctly updated. The numbers that get fed through to you could be incorrect. People should be prepared, on a cost-benefit basis, to spend time performing data validation. Finally, once you have calibrated your models, perform back-testing to check that the outputs are consistent with the inputs.

Chart 2 provides a verbal interpretation of the steps that you'd go through, from beginning to end. There are a few points to keep in mind. No one model fits all solutions. You have to learn to walk before you can run. Keeping it simple is probably the best thing. There always seems to be a tendency to go the more complex route, whereas the 80/20 rule seems to apply—80 percent of the benefit is picked up in 20 percent of the effort. You can certainly get a great pick-up initially modeling the additional components, but going too far into the details may not be worth it. Finally, always strive toward actionable results.

I now want to talk briefly about random number generation. It's often overlooked during the model development phase. The objective is simple enough—to produce numbers that are uniformly distributed between 0 and 1. You probably never thought that random numbers could be such an exciting field, but it is a fundamental building block of any stochastic simulation. People rely on random number generators within software, but often they have no idea how robust the random number generation process is. If you look on the Internet, you'll find many references for random number generators that are available. Many academic sites indicate tests that you should run your random number generator through. However, no random number generator will satisfy all tests. They will repeat sooner or later.

A good practice is for your company to adopt a standard random number generator that's used for all stochastic modeling purposes. Then as other stochastic models are developed, they can rely on the standard random number generator. Also, it can take some of the mystery out of trying to debug unusual results. I have seen cases where the random numbers don't repeat very often. What ends up happening is that you run 1,000 scenarios, and after 200 or 300 scenarios, the scenarios start to repeat or there's something in there that's cycling. You thought you ran 1,000 scenarios, but actually you ran 250 scenarios four times.

An economic scenario generator is another important part of a stochastic simulation that contains capital market-related guarantees. Here, again, it makes sense to have a common generator within your company. There are many to choose from. You need to determine whether you require an economic or a statistical model. Calibration is always an issue. A desirable characteristic to look for is that the scenario generator is an integrated model. In other words, it not only models equity returns, but also interest rate yield curves or fixed income returns as well as inflation and currency, all in an internally consistent fashion. A few models I've seen don't model the relationship between interest rates and equity, but I think that's changing. People now recognize that if you're going to model GMIBs, for example, you have to capture the joint interest rate and equity market return relationship in your model. Another desirable characteristic that you want to have in your scenario generator is a component approach. Instead of having to re-run the whole model for just one part, it's flexible enough so you can run only the pieces that you need.

I now plan to use a GMIB rider as an example, because it is a fairly simple product (in concept at least). With this product, a person deposits $100. There's a guaranteed account and a market value account. The guaranteed account accumulates at a "roll-up" rate. In this case I've assumed 5 percent per annum. The market value account is driven by the market value of the funds (accounts) that the person invests in.

Let's assume that the nature of the situation, which is not unlike companies in the industry right now, is that we have a block of contract holders with an overall guaranteed account value of $1.4 billion while the market value is equal to $1 billion. So the guaranteed income benefit option is "in-the-money." The options, due to assumed exercise constraints, are three to four years from being able to be exercised. A key thing to remember about the GMIB is that there's a dual impact. You have an interest rate guarantee tied with a mortality guarantee. On one hand, you have a conservative interest rate set fairly low, maybe 3 or 4 percent on the annuity. (It was conservative when it was priced, but unfortunately it's not any more.) On the other hand, you have very aggressive roll-up rates. You're guaranteeing 5 percent, but the only way the person can realize that is by taking the guaranteed annuity.

I generated equity returns using regime switching log-normal model with two regimes. Fixed income returns were modeled using a Cox-Ingersoll-Ross model. Historical correlations were used. The economic scenario generator was calibrated using maximum likelihood estimation. There are calibration issues that you have to worry about that include limited, often inconsistent, data. The Choleski decomposition methodology was used to generate correlated returns. If you don't use time series of equal length, your model may eventually fail because you won't properly satisfy the requirements for the Choleski decomposition. There are various "fixes" for some of these problems, but they can become complex.

One logical question is, how often do you recalibrate the economic scenario generator? If you are presenting results quarterly and are also recalibrating quarterly, then changes in the results will be driven by the recalibration as well as by changes in the market condition. You can actually be making the presentation of your results more complicated if you recalibrate more often.

In my example, I ran 1,000 scenarios using a monthly frequency and a 35-year projection horizon. Chart 3 is a scatter plot of the present value of GMIB cash flows as a function of the average interest rate per scenario. In looking at this chart, the benefit is certainly under water—that is, it's "in-the-money." You can see there's a limited upside, given the modeling assumptions that went in. Note that whether interest rates are high or low doesn't make too much difference.

Chart 4 shows the present value of GMIB cash flows as a function of the average equity return per scenario. It's the same scatter points, but plotted in a different X-Y plane. Here, you also see there's a limited upside. It should be, because it's the same results being plotted as in Chart 3; if it were otherwise there would be a problem. That is one of the difficulties with stochastic modeling. Sometimes you're dealing with so much data and while condensing it down, it's easy to make mistakes. Note that as the equity return decreases, from the upper right to the lower left, there are a number of scenarios where the end result could be classified as catastrophic in nature.

The conditional tail expectation (CTE) is a measure of downside risk. It's defined as the average of outcomes that exceed a specified percentile. In other words, if you want to calculate CTE (90 percent), you order the results from 1 to 100 (if you had 100 observations) and average the ten worst results. It's considered a more robust measure than a percentile. When you're dealing with embedded options where you have skewed distributions, the CTE picks up on the catastrophic events in the tail. In some cases, the CTE is modified. For example, if you have events where you don't allow the upside to be averaged in, you just floor it out at zero.

Selected CTE measures for our GMIB rider example are shown in Chart 5. By definition, CTE (0 percent) is equivalent to the average result. The results are presented as negative present values of the GMIB rider cash flows, so on average the present value is minus $43 million. Recall that this option is "in the money"—that is, the option has value to the contract holder. What was thought to be a free option a few years ago, or given away for virtually nothing, is now seemingly very expensive.

The percentile and CTE curves are plotted together in Chart 6. I originally also plotted a modified conditional expectation measure where I zeroed out the favorable events. Unfortunately, because this product was so under water, the modified conditional tail was almost identical to the true conditional tail, so it didn't really capture the point that I wanted to make. Usually, by zeroing out the positive events, the modified conditional tail will result in a larger number. However, as you move into the extreme tail, they eventually converge.

In Chart 6, the present value of GMIB cash flows is presented by percentile on the left side and by CTE on the right side. One observation is that the CTE measure looks considerably smoother. That's because you're averaging events, so there's more information contained in there. Another observation is that the percentile curves have various crossover points.

I performed some sensitivity testing on the model. The base equity return was assumed to equal 8 percent. I reset the model, recalibrated it such that the average equity return was equal to 6 percent, and reran the economic scenarios. Under a lower equity market return, from a percentile and a CTE perspective, the results have deteriorated. I also independently cut the lapse rate in half and the cost of the rider skyrocketed.

There are a number of benefits from performing sensitivity testing. Aside from quantifying an impact of changing assumptions, it is useful for validation of the model. If you sensitivity test an assumption and the results don't bear out what you expected, then you can do two things: you can conclude that it's insensitive, or you can go back and check it. I don't think enough time is spent on the model validation step. Sensitivity testing, if the model is correctly specified and the interaction correctly modeled, allows one to direct more attention to the assumptions to which the results are most sensitive. If you're wrestling with an assumption for a long period of time and your results aren't sensitive to that assumption, then you are wasting your time.

When performing stochastic modeling, the analysis-of-change step is very important. Often there are a number of parameter changes. It is useful to construct a build from one point in time to the other. Aside from getting a better understanding of the model, you're confirming that the model is responding in a way that you believe it ought to be. One is able to gain a better understanding of the dynamics underlying the stochastic simulation by performing sensitivity testing.

One method of presenting sensitivity testing results is shown in Charts 7 and 8. On the very left side is the base case. The line going across represents the CTE (70 percent) level for the base case. I picked CTE (70 percent) level merely for a reference point; it could have been any other CTE level. By looking at the results in this fashion, you can see how the various CTE levels change when you adjust the various assumptions.

Based upon the second column in Chart 8, changing the rider premium charge by 10 basis points (which may not be possible on existing business) doesn't make a big difference. If you're thinking about changing the pricing spread that you use in annuities at payout time by 10 basis points, it doesn't make a significant difference either. If you alter your mortality assumptions, in this situation, the GMIB rider is so far under water that you're just tinkering on the edges.

As far as the lapse rates are concerned, you can see that if by some miraculous chance we have double the lapse rate, the results are more favorable and compressed from a risk perspective. On the other hand, if the lapse rates get cut in half, for example if the contract holders suddenly understand that the GMIB rider contains a valuable benefit, then you can see that the risk gets significantly worse, especially when measured at the CTE (90 percent+) level. The discovery of embedded options is interesting because they can be discovered in one of two ways. Either you find them, or someone else will find them for you.

Normally the policyholder annuitization rate is a critical assumption; that is the take-up rate, or the rate at which they annuitize. This assumption probably would have been more sensitive if it had been more "at the money." Still, the outcome of increased annuitization is negative.

Sensitivity to the investment assumptions is presented in Chart 7. I've scaled the charts the same between the investment and the liability assumptions so you can make a comparative analysis. What does not show up when you look at the assumptions independently are the multiple variable interactions that occur in reality. For example, there's often a weak negative correlation between equity market returns and interest rates. When the equity market goes down, it's not good for this product, because the market value of the contract holder's account has fallen, while over time the contract holder's guaranteed account is steadily increasing at 5 percent per annum. When interest rates fall, then the current pricing rate for the annuity is lower and it's starting to hit the guaranteed rate, so that's not good. With the GMIB, when you start looking at the various jointly determined risk numbers, you quickly realize that you cannot just add the risks. For example, a minus 25 percent equity market shock combined with a minus 200 basis point shock in the interest rate market can't be taken in isolation and added. Due to the non-linear risk profile of the GMIB rider, one needs to test the income and equity return assumptions jointly.

One of the advantages of performing stochastic modeling is the potential for reserve and capital relief. In the Canadian environment, the use of stochastic approaches is favored over deterministic approaches. Stochastic techniques are also useful in quantifying risk and can be used for risk management and hedging purposes.

Keep in mind that no one model fits all. In our company, we've tried to adopt a standardized scenario generator for various stochastic modeling purposes so we can focus on the specification of the asset-liability model itself and the end results as opposed to re-building the scenario generation process. However, you still end up calibrating the model differently for Canadian versus U.S. requirements, or for whether it's a pricing exercise or a risk-management exercise.

In stochastic modeling, you want to cultivate best practices. You want to avoid "groupthink," where you have a group of people who think that what they're doing is absolutely correct. You end up with less communication going out and very poor judgment being made. It becomes not just a black box of a model, but also a black hole of a department.

Keep it simple; keep it practical. Don't use a sledgehammer to crack a walnut. I've heard many analogies describing stochastic modeling, but the point is that you don't want to "over-do" the situation. Any type of model that you put together should be suitable for the situation.

Focus on accuracy first and precision second. I tend to round the numbers so as not to overly imply any precision, because to some extent you're misleading yourself. Add complexity on a cost/benefit basis. Some people want to include everything in the model and then it takes forever for the model to run. You want to keep it simple.

You also want to perform reality checks along the way. Stand back and assess the situation. For example, take a look at the simulated output and generate summary statistics. Make sure that what is coming out was what you thought ought to be coming out. You might find a problem. On the other hand, you might find that you need significantly more scenarios than you ever thought to get a degree of confidence that you are comfortable with. There are formulas for calculating confidence intervals around the percentiles and around the CTE measures; they are quite useful. It takes a surprisingly large number of scenarios to reduce the confidence intervals.

Avoid the creation of black boxes. There's nothing worse than having a model that no one understands. If you do create a model that is on a common platform that people around the company use, that's a good thing, because at least there's a common understanding as to what the model is and what it does. One problem with some software languages is open code. Anyone can go in and change things. It can be hard when you're running Excel VBA code to lock it down.

There are other issues to wrestle with. Some models generate more volatility than others. Look at the various approaches when it comes to financial reporting. For example, if you're using a regime switching lognormal model with two regimes, you'll get much more volatility in the results over time than if you're using a simple log-normal process. The trouble is, when you're using a log-normal process, you're not getting thick enough tails, whereas the regime switching manages to satisfy the weaknesses of a simpler model. In short, there are always trade-offs to deal with. However, with a regime switching model, it may be more difficult to interpret the results.

For example, in a regime switching model with two regimes, you end up with six parameters: two means, two volatilities and two transition probabilities. You calibrate to your data using maximum likelihood estimation. You add in one more monthly return and your revised parameter estimates appear to be dramatically different. You ask yourself how you rationalize one parameter set from the other.

Given those six sets of parameters, you can calculate the expected return and the expected volatility. Now, you've distilled it down to a mean and a volatility, and it's much simpler to understand and compare.

There are other things to keep in mind when you're looking at calibration. I've observed when calibrating, especially in a regime switching lognormal model, if you add three months of data into a time series that's even thirty years long, and those data points happen to be very extreme movements, then the calibrated expected mean can significantly overshoot or undershoot the historical mean. This really came to light when we were performing some earnings-at-risk analysis. We added three monthly data points that were extreme events to our data. When we re-ran the model, the quantified risk level declined. People were asking how the risk could decline given that we have added three months of highly-volatile data and we only changed the calibration or the scenario generator. It turned out that the model in this case happened to undershoot the mean and volatility, and in the previous case it overshot the mean or volatility. So the results can be sensitive to the dynamics of the calibration procedure.

We also want to make sure that the fixed income returns are integrated with the equity returns in the model, especially if you're modeling embedded options that have joint interest rate and equity risks. The question is how to capture the correlations between the markets. Another question is how many scenarios to use. You never get a good answer on that, and I'm not going to give you a good answer. One of the common answers is that you just keep running more sets of 1,000 or 5,000 and so on until you see that you don't get a material difference in the end results between the sets of scenarios. Some people just give up after 10,000, based upon systems imitations.

Policyholder behavior is a critical element when assessing the risk embedded within GMIBs and GMDBs, but how do we model it? We don't have a material amount of data on policyholder behavior. Whatever data we do have can be highly dependent upon the company and the product design, so how to model policyholder behavior is a tough issue to wrestle with.

MR. HENRY M. MCMILLAN: (Pacific Life) I'm curious about your calibration. It appears that you're normally working with monthly data for stock prices and interest rates. Is that correct?

MR. HARASYM: For the example I presented, that is correct.

MR. MCMILLAN: How far back in time do you typically go to get the information that you're looking at? How many months?

MR. HARASYM: It depends on what we're doing. For example, if we're doing Canadian annual statement reserves for segregated fund products, we'll have to go back to the 1950s because the regulations tell us to do so. If we're performing a pricing exercise, we may go back only 10, 15 or 20 years. If we're doing an earnings-at-risk, we'll probably look at various combinations of 10, 15, 20 and 30 years and try to get a process that we're comfortable with. We try to settle on a calibration process so we don't keep changing our process each month or quarter.

I guess the answer to your question would depend upon the product. For example, if we were doing hedging of a segregated fund or a GMDB, if we were trying to do dynamic hedging and fitting data to it, then we've gone to daily information. That's where we found there were a lot of data problems with fund information that was being provided.

MR. MCMILLAN: So the answer is that it depends. But if you were using monthly data, you would probably always look for something in the area of 300 to 500 data observations when you're trying to estimate the parameters that you have here. If that's daily data, then you have 500 daily observations, et cetera?

MR. HARASYM: Yes, you certainly want to have a reasonable time series. There is one thing to keep in mind though if, for example, you're looking at interest rates. We often talk about stock prices, but if you're looking at calibrating for interest rates and there's a different interest rate regime that you're in or a different macro-economic environment, then it may make sense to truncate it somewhere.

MR. MCMILLAN: That's absolutely right. I was wondering how you possibly dealt with that when you were getting to the interest rate process because that's obviously changed. I would expect that with you doing international activity as well, you might be interested in foreign exchange movements, and then you have foreign exchange regimes that you have to deal with.

MR. HARASYM: Absolutely correct; that's a good point. A lot of it is gut feel. The interesting aspect about foreign exchange is that we do have to model it for earnings-at-risk purposes because we have operations around the world. A number of academic papers out there say a basic geometric Brownian motion with a few minor adjustments is as good as anything. Much to my surprise, when talking to people in the banking community, that's all they were doing too.

MR. WILSON: I think it's safe to say if you're looking at foreign exchange, it depends on the time frame over which you were intending to look. If you're looking at earnings-at-risk, which probably covers a one-year period and is totally random, a Brownian motion probably works rather well. But if you're looking at doing some sort of projection from a strategy motivation, which looks over a 10- or 15-year period, then there probably is a relationship between interest rates and currency changes. But over a short period of time there isn't one.

For example, we have a fairly large operation in the Philippines. The Philippines currency against either the U.S. currency or the Canadian currency, over the last fifty years, has basically done nothing except decline. That shows in the interest rates as well. The interest rates in the Philippines have tended to be about 500 or 600 basis points above U.S. rates at the short end and maybe 10 percent points above U.S. rates at the long end. But their currency has actually strengthened over the last year against the U.S. dollar, not driven by the fundamentals of the Philippines economy, but driven by the fact that the U.S. dollar has sunk.

MR. HARASYM: There is one point that I'd like to make about the calibration procedures. It's easy to think that you have achieved a global solution by setting up the model using Solver in Excel. However, in some cases if you change your starting conditions you'll get a different answer. You have to look at your maximum likelihood function, whether you're minimizing it or maximizing it and whether you have a better answer or not. You can end up with a local solution without realizing it. I've seen cases where the Solver routine is placed into VBA code and the process is repeated in order to increase the likelihood that they have found a "better" solution.

MR. WILSON: I have another example, which I'm stealing from Jeff Hancock, who presented it at the Appointed Actuaries Seminar in Toronto a couple of years ago. When the CIA did its study that was included as part of the American Academy of Actuaries' recommendation on risk-based capital for calibrating the regime switching log-normal model, Jeff looked at monthly data from the end of 1956 through the end of 1999. It was monthly, but the results differed depending on whether you picked the first business day of the month, the second business day of the month, third, et cetera. So if you picked the first of the month versus the last trading day of the month for every month, you got a different answer, and it was significantly different in terms of the calibration of the model. That's one reason you want to use something like a CTE.

Remember the fact that these are models. There's no model out there that says "airplane hits World Trade Center." Those just become data points that the model tries to say are stochastic and, of course, are not stochastic in the normal sense but are in a chaotic sense. On the interest rates, the oil price shock in the 1970s, which helped drive interest rates up and stock markets down, is certainly not something that just flows out of the stochastic model. The model is devoid of any economic and sociopolitical ramifications.

You can get significant regime change. Andrew Smith, one of the theoreticians at B.W. Deloitte over in the U.K., came up with a model that had a provision for a random shock in it. So you have this nice stochastic model, and then suddenly it would go haywire for a period, which is probably more realistic. For years, we've been doing dynamic capital adequacy testing as required under Canadian regulations. When I became an appointed actuary at the company in 1998, we started putting in what we referred to as the "Japan scenario" for financial condition reporting. Of course, many people thought this was a totally ridiculous scenario that should never be included; it's a waste of our time to do it because it will never happen here; we're not Japan. The scenario had the interest rates going down to

200 basis points and the stock market dropping 50 percent. Last time I looked, NASDAQ was down 80 percent and interest rates are below 2 percent at the short end of the spectrum.

I once asked some people what embedded options they had in their products. They were selling regular and ordinary everyday life insurance, and they said "None." I said, "Oh? What do you call a 4 percent minimum interest rate?" On Wall Street that's called an interest rate floor. What do you call a cash value? On Wall Street they call that a put option. But they hadn't considered pricing it because in their minds that's not an option. "We've always had these things; we've had them for 50 years." They're options.

Danico Life is the largest life insurance company in Denmark. Its dividend formula has a provision that it has stochastically determined as to what the option cost is by interest rate guarantee. For its business that has 4 percent guarantees, the company has a 40-odd basis point charge in the dividend formula. I think it was ING or Aegon, one of those two companies, that had a factor in its European dividend formulas for even a 2.5 percent guarantee that was about 10 basis points, and then something like 50 basis points for a 4 percent guarantee. You can study these things stochastically, but you have to take all the answers with a grain of salt.

MR. MICHEL HEBERT: (Swiss Re) You mentioned that the random number generators are not all random. Do you have a comment about the random number generator in Excel?

MR. HARASYM: That is an interesting question because it seems to depend upon what version of Excel you're running. If you went back to Excel 95, some people had found that with a certain seed number, you could get the random sequence of numbers to repeat after 300 or so random numbers. From this perspective, the random number generator is unsatisfactory. On the other hand, we've done some testing on the random number generator in Excel 2000 and the results were much more favorable. The formula that Excel uses can be found on Microsoft's Web site.

What we do in our in-house random number generator is somewhat of a random pick. We generate a series of random numbers and then pick one randomly. We throw out the rest and then start again. On that basis, we perform testing on the random number sequences. There are two considerations: does the same number repeat or does the same number come back in the same sequence? Of course, other statistical tests are relevant. I would look on the Internet. There are a number of algorithms to test random number generators that are in the public domain, as well as random number generators themselves. People have gone through and tested them, so a lot of the leg work is already done. But you should still do the tests yourself to confirm them.

The random number generator is a very important thing and very easily overlooked. It's not going to get you on the front page of your Web site for generating a good random number generator. On the other hand, you might end up in the news for a bad product if you mispriced it because you had a bad random number generator. There is not a lot of glory in the random numbers, but they're very important.

MR. JAMES A. DEMOPOLOS: (ING) We're talking about situations where the amount of capital requirements will be dependent on whether or not stochastic modeling is used. Could you comment on the extent to which stochastically modeled results are auditable?

MR. HARASYM: Do you mean internally or externally?

MR. DEMOPOLOS: I was thinking in terms of external, but any comments would be welcome.

MR. HARASYM: The unusual aspect of auditing stochastic modeling is that it's often done in black box form. But good practices would be that the models are documented. Were you looking for some of the thoughts I'd have on what to search for when auditing?

MR. DEMOPOLOS: I was thinking more in terms of the fact that, because some of the interactions between variables are highly subjective, you could come up with two qualified actuaries giving different answers for stochastic measurements of liabilities. Would that create a problem in an audit situation?

MR. WILSON: Stochastic approaches for doing actuarial liabilities in Canada have been there for the last four years, so that our reserves on segregated fund guarantees, GMIBs and GMDBs are indeed calculated with a stochastic model. The regulator has come up with a long paper on what you have to do to be able to use your own models for capital purposes. We use the same model for doing our reserves. You can audit it, but you're right, you're going to end up with two actuaries who might come up with different answers. But two actuaries might come up with a different assumption for mortality. The assumptions under U.S. GAAP, as well as under Canadian, are the property of management. They are not the property of the auditor per se. The auditor can scream and yell if he or she doesn't like your expense assumption, your lapse assumption or your mortality assumption, but they're management's assumptions. As long as you can show the auditor that you have gone through a process that makes sense, you should be okay.

For the purposes of getting capital credit in Canada, our models have had to be certified by the regulator, so the regulator actually has come in. The regulator doesn't typically have people with this expertise, so they hire experts to work with their people to see whether or not the model makes sense. Every time we change the model we have to go back and have it audited by the regulator.

That comes from the banking world, because the Basil agreement on capital requirements for banks has moved to use of own models for credit risk in banks. They have to go through the same process of having all of their models certified by the regulator. That would apply to American banks as well if they wanted to use own models. There generally is an ability to use a factor approach, but it' generally produces larger capital requirements. Use of own models generally will give you a better capital position than if you just use factors. That's the direction that Europe is going as well, in the insurance industry. The insurance industry is going to have to have models to set its capital requirements because the world is being recognized as a very stochastic place.

Chart 1

Is There Really A Starting and Ending Point? … No!

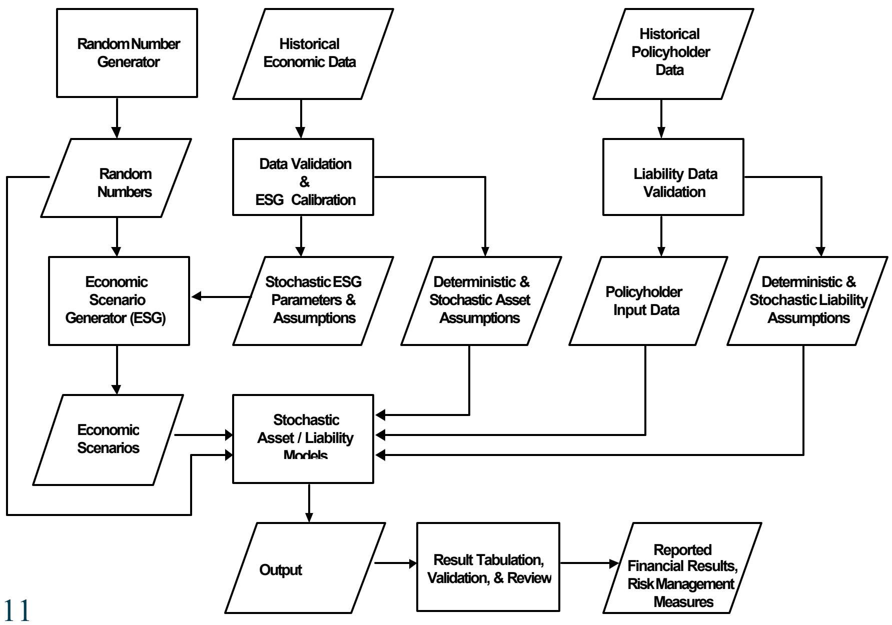

11

Chart 2

Where does one Start? Key Steps Are …

- Identify the key objectives and potential roadblocks before considering ways of solving the problem.
- Identify key issues and potential road blocks.
- Describe the process/model in general terms before proceeding to the specific.
- Develop the model: assumptions, input parameters, data, output.
- Fit the model: gather and analyze data, estimate input parameters
- Implement the model.
- Analyze and test sensitivity of the model results.
- Communicate the results.

12

Chart 3

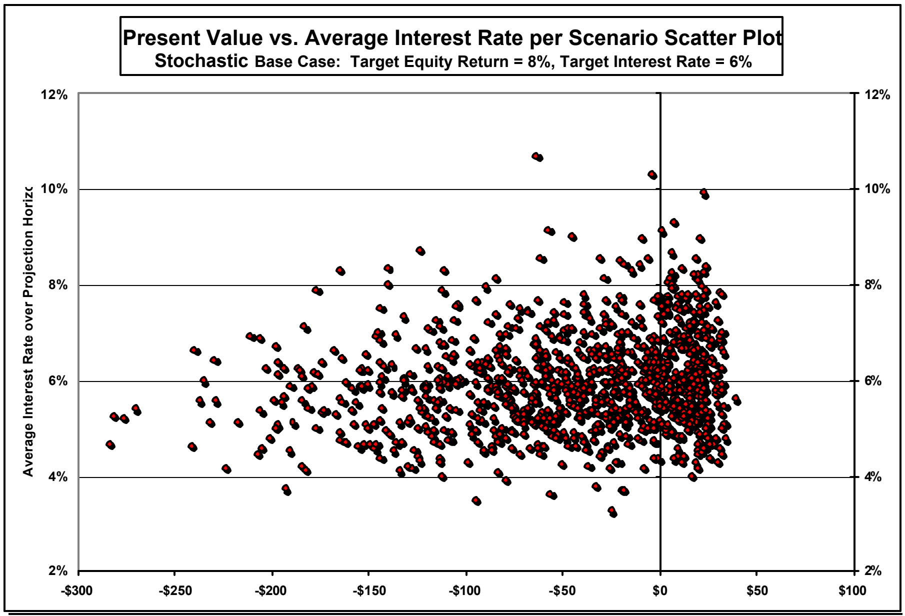

25

Chart 4

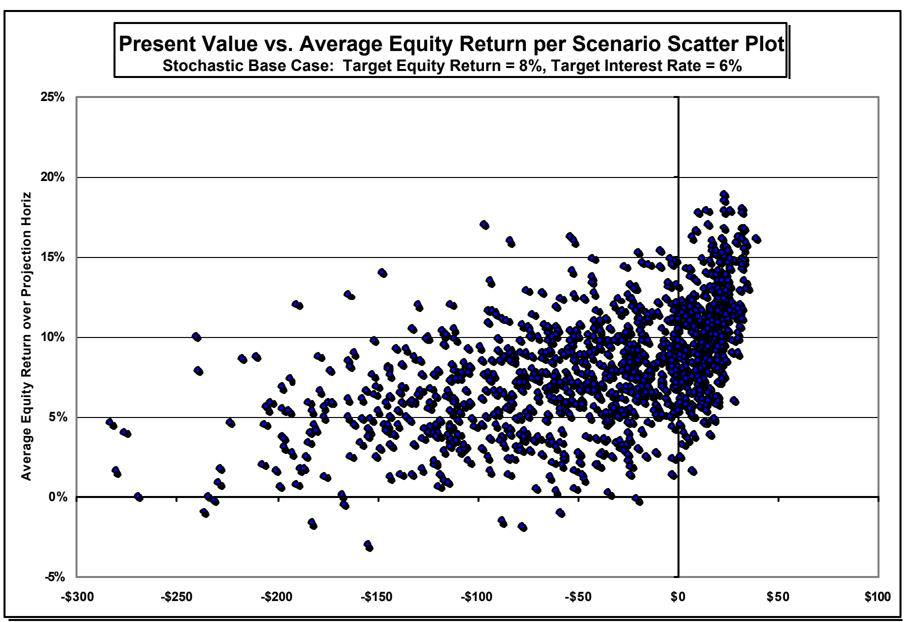

26

Chart 5

Stochastic Simulation Results:

CTE	GMIB ($millions)
95%	$204.3
90%	$177.2
80%	$145.8
75%	$133.9
70%	$123.8
65%	$114.9
60%	$106.9
0%	$43.4

- Recall
  - GMIB Guaranteed Account Value of $1.4B
  - Market Account Value of $1.0B

23

Chart 6

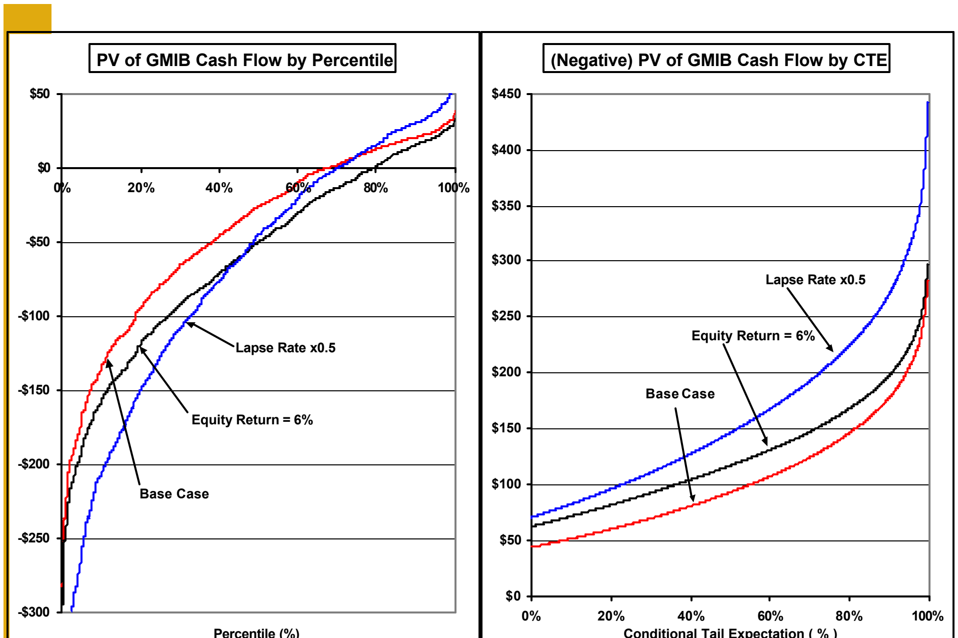

24

Chart 7

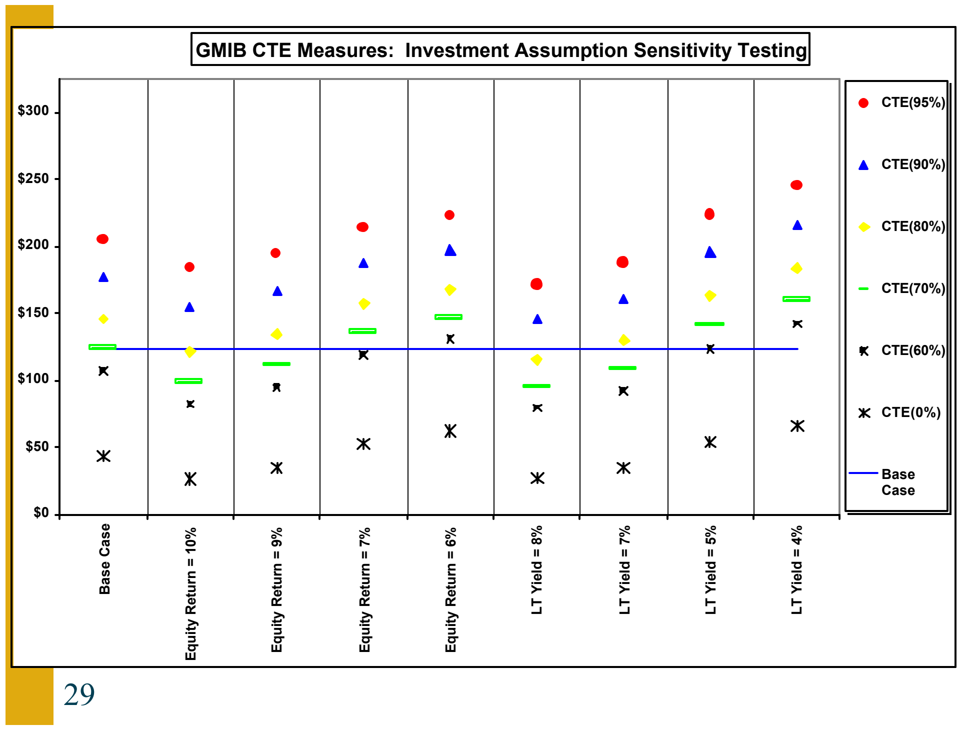

Chart 8