RECORD, Volume 25, No. 3*

San Francisco Annual Meeting
October 17-20, 1999

Session 130PD
Risks in Investment Accumulation Products: Recent Research

Track: Investment
Key Words: Asset/Liability Management, Equity-Indexed Products, Investment Products, Models, Valuation of Liabilities

Moderator: PETER D. TILLEY
Panelists: MARTIN K. LE ROUX
STEVEN P. MILLER
PETER D. TILLEY
Recorder: PETER D. TILLEY

Summary: The panel expands on research presented at the Actuarial Foundation/Nationwide Symposium held in January 1999. Topics include the potential costs of long-term guarantees, transfer risk modeling techniques, and the formulation of practical strategies to manage these products.

Mr. Peter D. Tilley: Steve Craighead was to be the moderator for this session. Unfortunately, he was unable to attend. We have made a few last minute adjustments. We have managed to put together the presentation as was originally intended. Steve gave a presentation at the Nationwide Symposium in New York City last January. I had the pleasure of serving on the Planning Committee with Steve, Irwin Vanderhoof, and several other distinguished actuaries. In an introduction to two days of paper presentations on risks in accumulation products, Steve presented a topic that's near and dear to his heart and, obviously, many others if we base our opinion on the attendance.

I'm responsible for the portfolio management of all of the different pension and insurance contracts that we have at Great West Life and Annuity.

Steve Miller is vice president of derivatives and portfolio manager at Mutual of Omaha. Steve is a chartered financial analyst as well as an FSA. He is responsible for analyzing and hedging interest rate risk for all the portfolios at all of the Mutual of Omaha Companies, and he has done some work on interest rate generators. His work has been published in *Risks and Rewards*, the newsletter of the Investment Section. He's a graduate of the University of Nebraska-Lincoln.

Note: The charts referred to in the text can be found at the end of the manuscript.

Martin le Roux is a senior investment officer at Sun Life of Canada. His primary activity is asset/liability management, including development and implementation of derivative strategies. Prior to joining Sun Life in 1990, he spent six years with the South African branch of a large U.K. life insurer. He's a graduate of the University of the Witwatersrand in Johannesburg, a Fellow of the Institute of Actuaries, and a Fellow of the Canadian Institute of Actuaries. He also holds a Certificate in Derivatives from the Institute of Actuaries.

We have a panel that should be able to enlighten you on some very interesting work that has been done in terms of long-term guarantees and segregated fund guarantees. I would like to recall the things that Steve Craighead said in his introductory presentation in January 1999.

Generally, Steve defined investment products, the topic of all the papers, as a contract designed to accumulate funds with the intention of making payments at some future date, usually, in terms of a retirement income-funding vehicle. It could be single pay; it could be annual pay. Annual pay presents its own set of risks. What kind of guarantees have you made for future deposits coming in? Examples of such contracts are mutual funds, bank CDs, and fixed and variable annuities. There are various tax considerations that might affect the pricing and marketing of these products.

The first set of papers dealt with economic scenario generation. In his introductory remarks, Steve characterized his own perspectives on generators. Stylized facts refer to terminology that you might be familiar with from some of David Becker's work that, again, has been published in *Risks and Rewards.* I think they are also in the monograph that came out from the Investment Section. Stylized facts are observances. If you look at U.S. Treasury rates, for example, over a large period of time, what kind of trends and patterns and what sort of things like frequency of inversions, for example, do you see? Does your generator fit these stylized facts or does it go off with its own mind and generate a set of scenarios that don't really fit anything we've ever seen in history?

Generators should be arbitrage-free. You should be able to use a generator to reproduce prices of instruments in the market. The logical structure of the generator should be consistent with respect to understanding the history of interest rates and macroeconomic behavior over long time periods, and I think that gets back to point one. It should model not just interest rates, but also things like inflation, stock market returns, dividend indices, and so on. In particular, if you've looked at things like equity-indexed annuities or the risks that are involved in mutual fund guarantees, segmented fund guarantees, then obviously you'd have to have some sort of a predictor or generator of stock market returns. The design of the generator should be modular.

The next set of papers was on behavioral microeconomics dealing with customer psychology. Steve's one comment on this is: are we thinking in a rational way or is rational thought not human at all? When we put our models together and we're doing things like policyholder behavior, (particularly lapses), we assume a certain efficiency of the use of policyholders' options to lapse their products. If my credited rate is 6%, and the competition's new money rate is now at 7%, then 28% of my policyholders might lapse, for example. Is that the way we can really think about our customers or are they a little more emotional in exercising these options than using rational left-brained kinds of thinking?

There were also papers on financial market risk, discontinuity, or catastrophe. There was a very interesting presentation by Mark Rubinstein of Cox-Ross-Rubinstein fame, who spoke of the market discontinuity. Twelve years ago it was Black Monday in the stock market, and what went wrong was portfolio insurance really blew up. He was involved and his firm was doing portfolio insurance for clients at the time and he maintained that he still believes in portfolio insurance.

I recommend the publication *Derivatives Strategy.*There was a recent interesting interview with Rubinstein on portfolio insurance. He believes that even though this shouldn't have ever happened, these kinds of things do happen. It's not necessarily a knock on portfolio insurance. The market lost all of its liquidity temporarily, and the portfolio insurance wasn't able to do the sorts of transactions that were required.

The members of the Planning Committee had hoped that there were going to be papers on the effect of distribution channels on policyholder behavior and persistency for this symposium. This was my personal contribution or lack of, I suppose, to the symposium, because each member of the Planning Committee was to come up with ideas and then there was a call for papers. I really wanted to see how the people in the industry and research would respond about the effect on persistency of selling through stockbrokers ??as opposed to a brokerage distribution system or a captive agent system.

We had described in the call for papers that we'd be interested in seeing some research on the effect of product design on customer persistency. If you are a student of the single premium deferred annuity (SPDA) market in the U.S., you've noticed that there are a few companies out there that offer a variety of SPDAs. The difference among their SPDAs, as best I can tell, seems to be that some of them pay a lower interest rate and have higher commissions, but the rest of the contract structures are the same. I guess I struggle with that concept of offering two different people two different contracts, where you could have given each person the same credited rate, but because you're paying a higher commission on one, you give them a lower credited rate.

There was an editorial in the *National Underwriter* recently; an agent was writing about one of these situations. The agent felt it wasn't right for a company to have five different versions of SPDAs. Essentially, the same contract is under five different credited rates with five different commission levels. The brokerage manager was trying to persuade him that this was great because he could get as much commission as he could make off a particular customer. The agent's response was, "Well, which one would you want me to sell to your mother?"

Again, I was hoping that we would see some research on whether there were different lapse rates on the ones that are sold with the lower credited rates and the higher commission or whether the agents are somehow manipulating customers this way. Unfortunately, there were no papers that came in. I had to respond with these sorts of comments to a fine paper that we commissioned from Shane Chalke, who wrote about the Internet as a distribution system for annuities. E-Annuity, which is affiliated with Lincoln National, is selling annuities over the Internet. Shane gave an interesting presentation on the customer psychology of that distribution system. He spoke in terms of what consumers expected for customer service, whether they ever really wanted to have any human contact, or whether they just wanted to get the entire thing done without ever having to speak to a live person. I'd recommend that paper as an interesting one to read when you get the book of papers.

Some comments that Steve had about distribution channel, again, would have been much more appropriate had we had lots of papers.

The distribution channel can have risks. You may be dealing with a distribution channel that is interested in getting as many commissions as possible from the same piece of business. That would be, for example, churning of policies.

There were also papers on long-term liabilities, options, and guarantees, which Steve Miller will be speaking about in more detail. We have to understand the economic conditions that create the exposure to the underlying risks. What kind of risks have we actually got on our balance sheet that we haven't really looked at very carefully? Think about some contracts that were sold 10–20 years ago. What kind of settlement option guarantees do they have? What does the risk on that amount do if you have some very low interest rate scenarios or even some mortality scenarios where the mortality improves dramatically and you have a long-term settlement option guarantee that is based on a higher death rate, an older mortality table? We have things like floor guarantees and guaranteed minimum interest rates on universal life and SPDAs.

Protecting fee income on variable products. We've written an option on our variable products that says, here's our mortality and expense charge, and this is what it will be for the lifetime of this contract. Are there any potential problems with that guarantee?

Transfer risk models is something that is very close to Steve's heart. He has done some research at Nationwide on transfer risks between fixed and variable funds. Think about a couple of examples of what can go wrong. Suppose you have a fixed fund where you're guaranteeing your interest rates for a deposit window of a quarter at a time. You have a particular credited rate that's out there for the fourth quarter of 1999. That credited rate is probably based on what interest-rate environment you were in, let's say, in the middle of September when you declared the rate for the fourth-quarter deposits. Suddenly, in the middle of November, the stock market crashes.

My own chief actuary at Great-West Life and I have arguments about this one all the time. He maintains that if the stock market crashes, people will flood into these growth funds knowing that this is the best possible time to get in. I maintain that

that's when people hit the panic button and they lock in their losses and they go whipping over into the fixed fund. If, at the same time, interest rates have gone down, you may find that the credited rate that you had set for these deposits can't be supported anymore.

There's another problem that Steve has identified where it's really a capital risk for you. Imagine your capital charge or capital allocation is ten times as high on the fixed funds as it is on the variable funds. Suddenly, a substantial portion of your variable funds transfer into fixed for some reason. What is this going to do to things like your risk-based capital ratio, or, for Canadians, the minimum capital and surplus requirements ratio?

There is a distinction between in-house and out-house transfers. An in-house transfer is one where your own funds are just moving back and forth. You have things like capital issues as I just described. An out-house transfer is one where you actually have competing funds alongside. We all try to avoid that situation when we're selling a case, but sometimes we're in situations where the customer can transfer to another company's funds and back and forth.

Disintermediation risk, of course, would be where you've sunk a number of expenses and perhaps longer term investments into a fixed fund. If it transfers over to somebody else's fixed fund on a different credited rate, maybe they moved because interest rates have moved up. Then you have some disintermediation risk on your own balance sheet.

Expense risk would be, again, a situation where you've sunk some costs in and now you don't have that block of business anymore to recover. Or perhaps the expenses that you're charging on certain funds aren't as high as they are on other funds. You had planned on a certain stream of expense revenues. There have been some fund transfers into, perhaps, the lower expense funds and you're not going to get that pattern of expense recoveries that you expected.

New money has lots of choices of places to go. You may have a variety of fixed funds and variable funds on an internal and external basis. I want to illustrate the point that if the customer has lots of choices that can be a very difficult thing for us to model.

With intergenerational equity, you want to make sure that you are not treating your old policyholders unfairly to attract new business.

For modeling considerations, Steve considered the Hippocratic oath here. First, "do no harm." You don't want to set up a model that ends up making you worse off than you were before you had the model. It has such bad assumptions in it or such inappropriate shortcuts that it actually steers you in a different direction for your investment policy or your crediting rate strategy. You have to have a good understanding of your competitor's credited rate strategy. As interest rates move up or down, what will the competitor do?

What I find challenging about this particular assumption is that it seems that your competitors are constantly changing. You may not be dealing with the same competitors in an up environment that you have in a down environment. In a down environment, it seems like your competition is the portfolio companies. In an up environment, your competition's the new money companies. It seems difficult to win some days.

You have to have a good understanding of your company's crediting strategy. I think the main point was that it shouldn't just be what you think the company's crediting strategy ought to be. It really has to be what your company is going to do in certain situations. You can't develop that kind of strategy sitting at your desk in the asset/liability management department. You must talk to the lines of business; you must talk to the marketing personnel, and you have to actually believe what they tell you and model that.

Then one of the things that you can do with this kind of information is show them the error of their ways after you've run it through the models. Perhaps, you'll get them to change their strategy. But if they insist that this is how they are going to handle various economic environments with the credited rates, then that's what you have to model. Otherwise, you'll end up driving yourself towards other kinds of strategies like investment strategies that aren't appropriate. You will, at some point, maybe not today, maybe not tomorrow, but soon, and for the rest of your life, regret this.

For product design, you have to accurately model the way your product works, current rates, guaranteed rates, surrender charges, the way that loans can be taken out, and the way that transfers can be done. Don't take any dangerous shortcuts just to make your modeling simpler. Don't avoid modeling particular options that the customer has.

We briefly touched on distribution channels. Two companies' SPDA models may look entirely different and for very good reasons. They might both be right for that particular company's distribution system. I might have a more sensitive policyholder behavior curve than someone else. That doesn't mean that they're wrong in their model or that I'm wrong in mine. It's just that if I'm selling through banks and you're selling through a career agency force, those are both appropriate assumptions. As with Occam's razor, usually the simplest solution is best.

There are several practical aspects of managing annuity blocks of business. This was a topic that Planning Committee member Timothy Pfeifer of Milliman & Robertson had brought up for papers. You have to understand your systems. You can't manage a block of business assuming you're going to be able to look at information daily if it's not available daily. If it's only available monthly or quarterly, you have to adapt your procedures accordingly. You have to have practical limits on how you manage your annuity block. You can't say that you're going to have a separate cell in your model for every single policy. You have to have some kind of reasonable way of grouping business so that can you can run your models in a timely fashion.

On the asset side, again, you have to understand the limits on information. You may wish to do all sorts of fancy modeling on things like asset-backed securities and collateralized mortgage obligations. But if you don't have the administrative systems to back up what you want to do, you're just going to have to live with those limits.

Setting the credited rate isn't something that you can do every day. There are administrative limits on how often you can communicate with the customers. I think most typically on portfolio annuities you'd be looking at changing the rate once per year. You have to build those schedules and the administrative lead times that are required to make changes into your plans for managing this annuity block.

Timing differences. I guess what comes to mind here is that you might have timing differences on the information that's coming into you and how fast you can react to it. You can get interest rate information on the market many times a day. But you have to deal with how often you can actually use that information and what you can do with it.

There are shortcomings for all types of modeling. Be very careful that you don't assume that your model is perfect and make very quick decisions based on the output of the model. Make sure you understand what's coming out of that model and why and whether you truly agree with the results. You have to remain competitive. That is a practical limitation that the marketing people remind us of on a weekly basis.

Mr. Steven P. Miller: I want to talk about interest rate risk on annuity purchase guarantees. It's something I haven't heard very many people talk about. I think that it may be a demographic problem that comes as people who have bought annuities start actually using them for retirement instead of using them as good tax planning for investing. The first thing I'm going to talk about is an introduction and the background of the problem. Then I'm going to introduce a concept called an arbitrage-free Markov chain, which is the framework with which we're going to explore the problem and, hopefully, give you some ideas on other possible problems where that might be a useful tool to solve things. We're going to then go into a very simple flexible premium annuity and discuss the results of that particular study. Finally, I will present an overview of some possible strategies for dealing with the problem.

Long-term interest guarantees. Deferred annuities are generally considered to be a relatively short-term product. They're of very short duration. Often, after the surrender charge period, you expect a large number of people to surrender. However, annuity purchase options contain very long-term guarantees. To a certain extent, when you sell an annuity to say a 45-year-old now, you may be making an interest rate guarantee for 50 years from now.

It has not been thought of as significant, and I will present information to show that, traditionally, it has been insignificant. When interest rates are relatively high, the value is very low. However, when interest rates are not relatively high, as in

1998 and early 1999, the industry's attention focused on the risk of low-interest rates.

The five-year U.S.-dollar swap rate, in many ways, is a better indicator of interest rates, especially for insurance companies and other financial companies, than the Treasury rate. The interest rates in 1999 got to a very low point. They hadn't been that low in 30-some years and, obviously, that made people start thinking about low interest rates. But it could have been worse. Although, the rates were low in the U.S. (according to U.S. standards), they were actually the highest interest rates among the four top currencies in the world.

What you see in Chart 1 is the U.S. dollar and the British pound, which, on intermediate term and long-term rates, was lower. Also included are the Euro and the yen. People often will say, "We'll ignore the fact that Japan has much lower interest rates because their culture is very different." Maybe there's a valid point to that, but I don't believe there's a valid point to saying that the culture of the Euro and the economies of the European Union are that much different than they are in the U.S. So there was certainly legitimate belief that interest rates were not as low as they could possibly go, nor that they could not get that low again.

I want to look at this particular problem on a strategic basis. One of the things that we've noticed at Mutual of Omaha is that we put an awful lot of work into looking at what a product looks like right now and on the day I'm pricing it or on the day I'm evaluating it. We also tend to forget that when we're doing that, we're going to be selling this product through all sorts of different interest-rate environments.

It's very expensive to have to refile products in lots of different states and to explain to agents why it is we're pulling their favorite product and giving them a new one. It is difficult to explain why we are doing all those kinds of things that would be perfectly normal in the derivatives market where a lot of these pricing models are used. But, in the insurance industry, we tend to work a little bit slower. One of the things that I wanted to look at in this particular study was, how is my annuity going to react through all sorts of different interest-rate environments? I wanted to use an arbitrage-free model.

When I use the words *arbitrage-free*," I almost always say that the term refers to the set of yield curves that are possible in your model. You can tell whether it's arbitrage-free by the fact that you can map that set of yield curves to a specific set of probabilities. However, people tend to get hung up on the probabilities and you can have an arbitrage-free model with realistic probabilities if you start with something that is arbitrage-free and just map it to different probabilities. That's not what I'm going to do here, but sometimes I just try to tell people what I really think arbitrage-free means.

I want to create a simple model that can be run quickly, because I want to be able to get lots of information about lots of different interest-rate environments. I don't want to have five minutes per run. I want to describe the risk under the widest possible interest-rate environments, and I want to be able to explore some of the contributing factors to the risk.

The first thing I'm going to do is talk about creating an arbitrage-free Markov chain, which, given one thing that I've given up, is not a very difficult thing to do. First of all, I want to describe all possible short-term interest rate states and you can describe them anyway you want. Essentially, those are going to each turn into one yield curve.

In my example, I said short-term interest rates can be from 0.5–26% and that I want to add that they are either generally rising, generally steady, or generally falling. I'll get to how I did that in a minute. Then I want to create a transition matrix among the states. What is the probability of getting from 6% short-term with a falling trend to 7% short-term with a rising trend in one period (in my case that was a three-month period)?

What I did is assume a certain amount of volatility. I used a normal curve and picked those points that were, in each case, representative of all the interest rates around them on the normal curve. I also assigned a probability to them and made a transition matrix among states. I said if an interest rate trend tends to be falling, there is a 60% chance that the next time it will still be falling. There is a 30% chance that it will change to steady. There is a 10% chance that it will rise. This was a relatively arbitrary way of looking at things.

I picked numbers that, in the end, made me feel comfortable about my stylized facts that we talked about earlier. As a matter of fact, in the case of that trend, if you say that falling interest rates tend to be inverted yield curves, you can, using some math and Markov chains, find out what percentage of the time your interest rates will be falling. I was a little bit more symmetric than that.

Now I have two different things. I have a probability based upon volatility of getting from one interest rate to another, and I have a trend. Regarding the trend, I say that instead of being on a rising interest rate if we have 5%, instead of centering it on 5%, I'm centering it on 5.5%. That's how I put my rising trend in there.

After I create this transition matrix, which is now 156 by 156, I can calculate the spot prices on the short-term interest rates. That's very easy for the first one because I'm using continuous rates, so it is just -2.5 times the continuous rate. The two-period price is going to be the present value of the present value of the expected value of the spot prices two periods from now. I have one- and two-period prices for every single one of my yield curves. Then my third-period price is going to be the present value of the expected value of the spot prices two periods from now. I can go out as long as I want.

I went out, I believe, about 40 years in this particular case. Now that I have a term structure in terms of spot prices, I can calculate my spot, my forward, and my current coupon rates for my spot prices.

Just as an example, I wrote a little program in Visual Basic, and I got this particular part of the process done. It took about five seconds on a Pentium II. Now I have my model completely specified and it's arbitrage-free. I know what the probabilities are that will make everything price out to the current state, because that's how I calculated the states i.e., what the present of value of going from one place to another is. The current yield curve is not in there. I essentially said that I'm going to model the universe so that every yield curve possible is one of these 156 yield curves.

To make it a little easier to see, Chart 2 is an example of when you have a short-term rate of 5%. This is what three different yield curves look like. For a steady trend, it's a slightly positive interest rate. I threw a little mean reversion in there, so if it was at 25%, it had a tendency to go down. The rising trend tends to be a steeply positive yield curve, although, it does come down a little bit towards the end. The falling trend tends to be an inverted yield curve. You could certainly say that there are many yield curves that could have a short-term rate of 5% that don't look like this. At least we have the basic three: in this case, a steeply positive, an inverted one, and a flat yield curve.

The advantages of a Markov chain model are: it gives you a wider variety of yield curve shapes than a lattice, and it has a natural boundary for the interest rates. In my model, I said that there is no probability of going higher than 26%. There isn't a problem of interest rates going way too high or low. Also, backsolving, as you would use through a lattice, actually gives you the values at all states at once.

The disadvantages are that the yield curve shapes are certainly not as robust as they are in continuous state models. When you get towards the boundaries, they have different shapes, but they're all either inverted or positive. The biggest disadvantage of this is that the current yield curve is not represented inside in the universe. You can't use this to go off and price today's derivatives unless you somehow improve the model so that you can get the current yield curve as one of the states.

But right now this is not useful for that. This is useful for looking over a large possible number of yield curves. For path dependent problems, they require Monte Carlo methods. If you're using Monte Carlo methods, then there probably are other better yield curve generators to use that will solve some of these other problems.

We're going to actually work on some back-solving techniques, even though it's sometimes difficult to find a back-solving kind of path independent problem in an insurance world. In backsolving you start at the back of the projection. I assume that everybody's going to die at 115, so I start at age 115. Then I back up one period and I look at all possible things that could possibly happen one period forward, and I calculate the expected present value of the cash flow at age 115. Then I back up, again, and I take the cash flow at all possible places I can go to from each node, and I calculate the present value of that cash flow, plus the present value of the present value that I calculated in the previous statement. Then I do that for each node and I repeat until I get back to the current age.

Unlike a lattice, I start at the end. I have fewer on the back end, and far more on the front end. I end up with the answer under every possible interest rate scenario. That is the kind of information that I'm looking for, for this strategic kind of question: how is my product going to react under all sorts of interest rate scenarios?

However, the big thing is that it has to be path independent. You have to be able to stand at a node and know everything that could possibly happen and not have to know about what happened in the past. Because I'm standing there, the only things I know when I'm doing this problem are the state of interest rates and the time.

Cash-flow functions that depend upon history can't be solved using this. This actually disqualifies most crediting rate strategies, because most crediting rate strategies may involve the yield on a portfolio.

For example, if you have a relatively simple mortgage model, you can solve the value of a mortgage or a mortgage pass-through as long as your simple mortgage model only says that the prepayment depends upon interest rate and seasoning.

Now we're going into the point where we're just going to just try this out and look at a simple flexible premium annuity. I'm going to say let's issue to a 45-year-old female. It turns out that I could probably fix that pretty easily and with the results that I'm going to get, I could probably value a whole group of issues. It's going to pay a $250 premium every quarter, $1,000 a year. Withdrawals are, in this case, not going to be interest rate sensitive. They can be 5% before age 65, 25% after age 65, but 100% at age 80.

The premium persistency I'm going to use is 90%. We're going to change that and see how that affects the value of the guarantee that I'm talking about. I have included a mortality table. I'm going to guarantee a credited rate of 4% and will have a guaranteed purchase rate at 4%. By a guaranteed purchase rate at 4%, I mean that I'm going to use 4% with the annuity 2000 table, and that's going to be my guaranteed purchase option rate. I'm going to ignore expenses. It's always much easier to say I'm ignoring expenses and taxes.

We will start backsolving through the chain, and we're actually going to do it in three steps, because what I'm trying to do is get some useful per-dollar things that are path independent. The first thing that's going to be relatively easy is I'm going to backsolve through the chain and figure out the value at each possible place I could be—a combination of interest rate and time. I'm going to figure out the value of my purchase guarantee per dollar of account value.

I'll take my mortality discount, all of the possible cash flows at the spot rates, and the guaranteed purchase price. We will express that as a percentage of the life annuity before the guarantee. I will calculate this for every interest rate and every age. Another thing that we could do is assume in real life that there are some people who won't take a life annuity no matter how great a deal it is. You could say we can modify the value of this particular option at this step to reflect exercise inefficiency.

I have the value of the option that I'm looking for, and what I want to do is calculate the value of $1 of account value, which is something else that I can know at a state. I don't necessarily have to know the history of the policy to know how much account value it has. It's going to be the expected present value of one plus the credited rate for a period. That means that I have to know that my credited rate can't be path dependent. I want to adjust withdrawals, and multiply by the guaranteed percentage that I just calculated.

I'm going to assume that if somebody is making a withdrawal and the annuity purchase benefit gives them a better deal, they'll do that.

For the people who continue on, I'm going to credit them interest, but then I'm going to have one minus the withdrawal rate minus the mortality rate in my Qx^t. I'm going to take the value of the guarantee for the account value in the next period and go backwards. Let's say it turns out to be $1. In the next period, I'm going to go backwards and pick up that $1 and calculate its present value for my survivors. I'm going to calculate this value from back to front. At the beginning, I have, for any particular account value, excluding premiums, what the account value would give me if nobody paid any more premiums, and what the benefit to those people would be that I'm guaranteeing them.

From the Floor: What are you using for a discount rate?

Mr. Miller: I'm using the short-term interest rates coming off of the scenario. The third step is to value the guarantee and the value of $1 of annual premium. Once again, I can look at somebody and ask how much premium they are paying. I don't know have to know how much premium they've paid in the past. If I know that, then that's path independent and I can apply my premium persistency to that in the future; it's the same thing.

I go through and I say the value of $1 of premium creates $1 of account value. It's the value of my guarantee of $1 of account value, plus my premium persistency times the expected value of my future premiums. Calculate this value from back to front. The result at time zero is the value of a guarantee of new issue, the first time somebody gives me a premium.

In Chart 3, I'm going to make some changes. I'm always limited to only being able to do path-independent changes in this particular kind of case. If you don't do that, then you can use the Monte Carlo method. But, like I said, you might as well use some other model. The first thing I'm going to do is always credit the guaranteed rate.

For example, we discover that if you do that, then unless interest rates are very, very low, you don't have a huge amount of risk. If you look at the range from 5.5% to 6.5%, you find that you will probably have 5% of one year's worth of premium as the value of this particular option. Obviously, that's an unrealistic credited rate, but at least now I have some sort of idea of what the values could be.

I'm going to go to the other extreme in Chart 4. I always credit the short-term interest rate. I could have picked any particular interest rate, but it has to be a path-independent type of thing. I go from always crediting the guaranteed rate to always crediting the new money rate. I try to examine, at least on those two extremes, any difference due to the crediting strategy in the value of this particular option and I get almost exactly the same graph. As a matter of fact, I had to look individually on the pieces to make sure that I hadn't somehow accidentally run the same model, but those numbers are different than the previous numbers.

This makes me feel good. These particular models do not have any very sophisticated crediting strategies. At least given two opposite ends of the spectrum, it doesn't make very much difference. From now on I will always credit the short-term interest rate.

Now I'm going to look at something else that seems like it would make a lot of difference—premium persistency. For Chart 5, I'm going to use 100% premium persistency. It turns out that it's going to be important, but I also have to remember that I'm only looking at one aspect of profit or risk. A 100% premium persistency is, generally, a very good thing and so the excess profits may offset these things quite handily.

If I look at 100% premium persistency, I feel that it is something that could be relatively important. I looked at Chart 5 and noticed that it's almost like paying a 10% extra commission. I think that that might be something that I really want to deal with by having a different mortality assumption. Remember, this is a 4% purchase guarantee and a lot of people might use 3% purchase guarantees. This is something that I would look at as something that could be an important consideration if I have good premium persistency. But, once again, I have to look at my offset from profits. Could I pay a 10% extra commission, for example, if I had 100% premium persistency?

The last example is Chart 6. I was concerned that I would have anti-selective premium persistency in this chart. This is a flexible premium annuity. People can give me more. They don't always have to give me the same or less and so I have a simple rule. I'll assume 90% premium persistency if the credited rate is above 4%. But if the credited rate is equal to 4%, which is the minimum, then I'm going to make it 110% premium persistency, so that some people can be doubling up. At extremely low interest rates, this is much worse. However, the effect does tend to tail off when you get interest rates that are higher. This is the reason why the charts have been so big. I could fit this one in.

I've actually started identifying something that I have to be careful with. The worst possible case seems to be the most logical one, which is that people stop paying me premiums and go off and do better things unless they have a really good deal by paying me premiums and then they start being selective. Then I discover that I even have some comparisons, which could be important options.

Looking at one particular yield curve, I find out that as a percentage of premium, this is worth approximately 4% of $1,000 or $40 per $1,000 with the guaranteed and maybe a little bit more at the new money credited rate. It's worth nearly $95 per $1,000 at both the 100% premium persistency and at the antiselective premium persistency. The antiselective premium persistency value grows a lot as interest rates are going down.

Given this kind of information, we can start thinking of possible strategies. One of the things we could change is the product design by including some sort of limited flexibility or some sort of product design to help limit some bad behavior.

One thing that has been suggested in our company is to sell the current product, only if interest rates are above a certain level, but like I said before, there are a few problems with logistics. That's probably not a useful type of thing. If you happen to be a company that has the ability to do that, you could have a level at the very beginning and say that you are going to sell this product with these features under these circumstances, but if that changes, you are going to have a different product for people to buy. Hedge it with floors and call swaptions.

A final idea, and one I think that wouldn't be a bad idea, is to accept the risk. Look at this kind of information and adjust your target surplus and capital requirements so that if you decided that it got to be too much of a risk, you now had some surplus that you could use to hedge the remaining risk.

Mr. Martin K. le Roux: I want to speak about some work that we've done in applying option pricing models to analyze what are known in Canada as segregated fund guarantees.

To start, I should explain some of the terminology that is used in Canada. A Canadian segregated fund contract is basically the same as a separate account variable annuity contract. In other words, it's a mutual fund that has been dressed up to look like an insurance contract. In order to make it a valid insurance contract, you have to have some guarantees that go with it. Traditionally, in Canada, these guarantees have been quite modest, but about 3–4 years ago, a number of companies started to offer guarantees that were fairly significant. A typical arrangement might be that you would guarantee 100% of the initial deposit at contract maturity after ten years or on earlier death.

Some companies, although not quite as many these days, have also been offering what is known as a reset feature. This allows the contractholder to ratchet up the guarantee level to be equal to the current fund unit price. If you do that, then what happens is that the contract maturity date is pushed out, so that it's then ten years from the date at which reset takes place. In effect, this gives the contractholder the right to lapse and re-enter without any penalties in the form of additional sales charges or fresh underwriting, and so on.

Fees for these guarantees have typically been about 50-75 basis points per annum greater than an equivalent contract without the same level of guarantee. Total fees are typically about 50–300 basis points per annum.

These contracts have been widely successful from a marketing point of view. They've been enormously appealing to GIC refugees, people who are looking for an alternative to GICs in today's low-interest-rate environment, but are nervous about stock market volatility. It has created quite a challenge for the regulator and the actuarial profession in Canada. The existing reserving guideline is widely recognized to be no more than an interim stop-gap measure, but there is very little agreement on what ought to replace it.

The Canadian Institute of Actuaries published a discussion paper on this topic in 1998. The Institute handed the issue over to the profession for a debate. That debate took the form of a symposium in Toronto. There were no fewer than 26 presentations in a two-day period, which I think must be a record of some sort. This presentation was one of those 26. But it's not clear to me that much in the way of consensus emerged from that symposium.

One of the major issues is the question of whether one should be reserving for these contracts, using what I will call a statistical approach, sometimes known as an actuarial approach (and I'll explain in a moment what I mean by that) or whether one should be using a risk-neutral option pricing approach or some combination of those approaches.

A secondary issue is, are you are going to be using an option-pricing approach? Is something relatively simple, such as the classic Black-Scholes option pricing framework adequate given the long-term nature and complex nature of these guarantees or should you be seeing something more sophisticated than that? There's a quotation from Einstein: "Everything should be kept as simple as possible, but not simpler." The question is whether Black-Scholes meets or fails that test in this context.

To start with a description of what I mean by the statistical approach, the idea here is you start with a stochastic model of how stock prices or fund unit prices can evolve, and that could be something relatively straightforward such as a constant volatility lognormal process. You could use a variation on a lognormal process. There's something called a Regime Switching Lognormal Model that I'll have a bit to say about later. You could use a Wilkie Investment Model. There are many potential choices that you can use.

Once you've got your model of stock prices, you can, in turn, use that to model your liability cash flows. You can determine the value of the assets that you would need to support those liability cash flows. The way that you come up with a reserve is that you take a confidence limit of the resulting distribution. The Canadian Institute of Actuaries has suggested that since these guarantees are quite risky, if you're going to be using this approach, then probably you should be setting your reserve at something like an 85–90% confidence limit.

There are some issues with this approach: theoretical and practical. Here's the theoretical issue. Suppose that you've written a contract where your liability is just the stock price and that you've promised to pay on death or on contract maturity an amount that is equal to the value of some external mutual fund. From the point of view of the contractholder, if they own one of these contracts, they might as well go out and buy the mutual fund. The two are equivalent in economic terms.

But if you apply the statistical valuation technique projecting out the stock price, discounting back at some discount rate, taking a confidence limit of the resulting distribution, you will come up with a valuation that is probably quite different from the known market value of that stock. From a financial economics perspective, that would be considered a major problem. A financial economist would ask, if this technique doesn't even reproduce the current known value of a stock, how can you trust it to value an option on that stock which is what these maturity guarantees really are? Actuarial opinion seems quite divided on this point. Some actuaries side with the financial economists, and others claim that there really isn't any issue at all.

Regardless of where you stand on the theoretical issue, there are, however, a number of practical issues in using this approach. The first is one of computational inefficiency. You basically have to implement this using Monte Carlo simulation. Although there are ways in which you can speed up Monte Carlo, it is probably not the methodology of choice when you're dealing with many tens of thousands of individual contracts on a production basis.

These problems get far worse if your investment strategy is not simply a matter of backing your reserves with a passive portfolio of bonds. But if you're trying to hedge against the guarantees using some sort of dynamic rebalancing strategy, then the computational problem becomes much worse. Because what you have to do at every single time step for every single simulation is model the composition of what your hedging portfolio looked like. That multiplies the computational load by a factor of 100 or 1,000 or worse. You're talking about maybe 10 million computations for each individual contract. That is clearly not feasible.

The other issue that you run into is one of parameter uncertainty and parameter estimation. Even if you use something relatively straightforward such as a constant volatility lognormal model, you still have at least three subjective parameters that you have to estimate. You have to decide on an appropriate confidence limit, an expected return parameter, and a volatility parameter.

What you will typically find for these long-term, deep, out-of-the-money contracts or options is that the results you get are hugely sensitive to very small changes in those parameters. Even if two actuaries were to agree on the basic methodology and agree on the model to use, and even if they have slight differences in the premium as to what appropriate parameters might be, they will still come out with hugely different results when they use this technique.
In contrast to a statistical technique, one has the option pricing approach or, more specifically, the Black-Scholes option pricing framework. Now, the starting point here is quite similar. You're starting, again, with a stochastic model of the stock price evolution. In the case of Black-Scholes, you're using a constant volatility lognormal process. Using that model you can, in turn, model what your liability cash flows are. But the way in which those stochastic liability cash flows are translated back into a single net present value is fundamentally different.

Instead of using a statistical confidence limit, what you instead look at is the value of a hedging portfolio. This is the fundamental insight behind Black-Scholes: provided you make certain idealized assumptions, it turns out that you can construct, using a dynamic rebalancing strategy, a perfect risk-free hedge against your liabilities no matter what happens to the underlying stock price or fund unit price. You can be perfectly hedged against that. If you can do that, which at least in principle is possible, then the economic value of your liability is simply equal to the loan value of your hedging portfolio.

The enormous appeal of this approach is that it cuts right through the problems of parameter estimation and parameter uncertainty. You don't have to decide on a confidence limit because the results are the same in every single case. You don't have to decide on expected returns, because no matter what you expect returns to be, you end up being hedged against them. You do have to decide on a volatility parameter, but the sensitivity to that parameter is not nearly as severe as is the case for the statistical method.

The particular model that I've implemented is based on a binomial tree methodology. I've written a background paper, by the way, that describes the technicalities. Binomial tree framework is a pretty well-known methodology. Essentially, what you're doing, instead of working directly with a continuous time lognormal process, is making an approximation by using a discrete time binomial process. What that means is if you map out all possible paths for the evolution of your underlying stock price or fund unit price, you end up with something that looks like Chart 7. Chart 7 is an example of a binomial tree.

The way that you use this to value your liability is you start at the terminal nodes for the tree, the maturity date of the contract. At each of these nodes, you know what the stock price is, so you can determine what your guarantee is worth at that point. Then you take a final pair of nodes in turn and you work backwards to their common predecessor node. What you then do is determine the value of your hedging portfolio, in other words, the value of your liability at that predecessor node. You repeat this process for each pair of nodes, and gradually work backwards until you get to the origin of the tree, which gives you the present value of your liability.

This is a very powerful and versatile technique and, in particular, you can extend it quite easily to deal with quite a few features of segregated fund contracts. Just briefly, you can incorporate fund management fees at each node. It's very easy to determine what your fund management fees would be at that point because you know what the stock price is and you simply net them against the value of your liability. As an aside, that means that your liability could become negative. If the stock price is high, the present value of the guarantee is probably quite low. The fee income is probably quite high, so that could lead to a negative value.

You can also extend it to deal with mortality, lapses, and withdrawals provided that you're prepared to model these in a deterministic fashion. If that is the case, then basically what you do is overlay a traditional actuarial decrement table on top of the basic tree structure. The technical details are in the background paper.

I would just add that deterministic does not mean invariant. You could, for instance, have a lapse rate that is a function of the stock price. You might want to assume that if the stock prices plunge, then you will have an up tick in your lapse rates, and you can do this using a binomial tree approach. The only thing is it's a deterministic relationship, the lapse rate is not a stochastic variable in its own right.

It also gives you a very powerful and elegant way of dealing with resets and antiselective lapses. Reset, you'll recall, is basically just an antiselective lapse and reentry. The key assumption here is that your replacement contract is fairly priced and that the value of your liability at the point of reentry is precisely zero. That may be a fairly strong assumption, but if that is the case, then it follows that perfectly optimal withdrawal behavior or reset behavior is for the contractholder to monitor the value of their contract.

If it ever becomes negative, meaning that the present value of future fees is greater than the present value of future benefits, then, at that point, it's theoretically optimal for them to exercise their reset option. The way that you incorporate this in the model is as you work backwards through the tree, you simply determine whether your value at any node is negative and, if it is, you assume the reset takes place and you bring it back up to zero.

That's a fairly simple model. It is probably unrealistically conservative because the reality is that most of your contractholders don't have access to option valuation models. They're not going to be checking stock prices daily to see whether they should be resetting. A more realistic alternative might be to assume that if the current contract value goes negative, then some people will leap in and reset, but some fixed percentage of the population will simply allow their contracts to remain in force. Either way, one of the big advantages of a binomial tree approach or a trinomial tree, which is an alternative that I also use, is that it allows you to model optimal reset behavior in this fashion. It's not impossible with Monte Carlo approaches, but it's certainly a lot more complex and difficult.

The big advantage of the binomial tree is sheer computational efficiency. Provided that you're dealing with a single stochastic risk factor, this approach is easily 50-100 times faster than Monte Carlo simulation. A particular implementation that I've developed, which is put together in Visual Basic for Excel, which is not exactly the world's fastest programming language, ticks along quite nicely at about one second per contract, which is that ten-year contract working in 120 monthly time steps. It's quite feasible to implement this on a production basis.

What I now want to show you are some charts that compare these two approaches, the statistical and the option-valuation. Chart 8 is meant to illustrate this problem of parameter sensitivity. I'm starting with the statistical model. This is a specimen contract, face value of $1,000 currently at the money, which has five years to run to maturity. On the horizontal axis, I have the volatility parameter. On the vertical axis, I have the value of the liability. The different lines correspond to different values of the expected return parameter.

What you can see, for example, is if you were to change your expected return parameter from 12% to 10%, that would make about a $75 difference to the value of this contract. That's significant when you consider that your annual management fee or the portion of the fee that's available to fund the guarantee is only $7.50. Similarly, you can see that if you change your volatility parameter from 15% to 17%, then, again, that makes about a $75 difference to the value that you place on this contract. If you contrast that to the option pricing price approach, there is, of course, zero sensitivity to the expected return assumption. There is some sensitivity to the volatility assumption, but it's not nearly as severe.

On Chart 9, I've tried to show how these different models behave in response to a change in the current fund unit price. Again, starting with the statistical model, I'm keeping the expected returns and volatility parameters unchanged, but I've also shown what happens if you vary the confidence limit. Again, if you change from say a 90% confidence limit to a 95% confidence limit, that makes about a $100 difference in the value that you place on this contract.

The other thing you can see is that as the fund unit price moves up and down, provided the fund unit price is in excess of a certain threshold, there's basically zero sensitivity to a fluctuation in fund prices or stock prices. As soon as stock prices fall below that threshold, the values that you get go rocketing off into the stratosphere, which is implausible from a general economic perspective. It is also not terribly helpful from the viewpoint of a CFO who's trying to manage the emergence of quarterly earnings. If you contrast that with the option-pricing approach, instead of having this curve with a sharp kink in it, you get a nice smooth curve with a comparatively gentle gradient. That's my sales pitch for the option pricing approach, but there are some shortcomings as well.

The single biggest shortcoming is your reliance on the constant volatility lognormal model. Although, this is a pretty good assumption most of the time, it's well known that it underestimates the probability of extreme price movements relative to what you actually see in real life markets. The one example everyone brings up is the crash of 1987, which, if you believe in a constant volatility lognormal model, ought not to have occurred in the entire known history of the universe.

Other issues are the model assumes that interest rates are constant. That is clearly not the case in reality. It assumes in order for hedging to work perfectly, you have to be able to rebalance at infinitesimally small intervals with zero transaction costs, which is clearly unrealistic.

The final problem with segregated funds is that it also assumes that you can take a short position in mutual fund units, which is also pretty problematic. In reality, if you were to try to hedge these things, you would likely end up using index futures, but that introduces a fairly significant basis risk.

There are quite a few extensions to Black-Scholes that are designed to address these issues. As a general rule, what you find once you relax the constant volatility lognormal assumption is that perfect dynamic hedging is no longer possible. It's no longer possible to create a perfect risk-free hedge.

At this point, some people are inclined to throw their hands up and claim that option pricing is really completely useless. That's quite a strong reaction. My own viewpoint would be that while you may not be able to hedge these risks perfectly, you can still get a pretty good idea what they are worth by looking at the value of a hedging portfolio, even though your hedging strategy might be imperfect.

I try to illustrate this point by seeing how well a hedging strategy based on the plain-vanilla Black-Scholes approach would work under some conditions that are a little bit closer to the real world. I've used something called a Regime Switching Model. I would like to explain what that is, but I'm running a bit short on time. I've thrown in as many other real world restrictions that I could think of. I've assumed that portfolio rebalancing only takes place monthly. I've allowed for transaction costs. I've allowed for imperfect correlation between your underlying fund unit price and your hedging instrument.

I would add, by the way, that this is a very simple unsophisticated hedging strategy. In reality, if you were going to try and hedge these things, you would presumably be doing something a little more robust and sophisticated.

Here are my results. First, these are based on 10,000 simulations with a contract with a face amount of $1,000. If you leave the risks unhedged, about 55% of the time, the insurer would make a small profit between $25 and $50 per contract. That's just their management fee income over the course of the contract. Another 25% of the time the insurer makes a profit between $0 and $25. But 20% of the time, the insurer makes a loss anywhere from $1 to $500.

The statistical approach to reserving really doesn't work terribly well. What you find is that the best you can say is that your reserves should be set somewhere out there in the tail of the distribution, but precisely where it is, is going to depend crucially on your particular choice of confidence limits and your particular choice of model parameters.

If hedging worked perfectly, you would find in each and every case of these 10,000 simulations that you would get exactly the same answer. That is clearly not the case in reality. In reality, you do get the distribution of outcomes. But by looking at that distribution, you still get a much clearer idea of what this contract is really worth.

The mean of that distribution happens to be very close to the plain vanilla Black-Scholes option pricing model. You'll find that compared with the unhedged distribution, the hedged distribution has a standard deviation that's about half. The tail-end confidence limits are much more closely spaced. Just as the case with the Black-Scholes model, there's much less sensitivity to model parameters. There is almost zero sensitivity to expected returns and comparatively low sensitivity to the volatility parameter.

All in all, the conclusion that I draw from this is that the basic Black-Scholes model, despite the idealized and simplified nature of its assumptions, turns out to be a pretty good guide to the fair value of this contract. However, in Canada we don't reserve on a fair-value basis. The reserving standard requires you to hold the best estimate, plus a provision for deviations.

You do need to have some additional provisions and that is where a statistical approach might well come into it. There are two reasons why you need these provisions. First, as we've just seen, you can't ignore the effects of model risks and parameter uncertainty altogether. Second is that you might very well not be hedging against these guarantees. You may simply be accumulating all your reserves passively in bonds. I stress that point.

What I am suggesting is that the way to value these guarantees is to look at the value of a hypothetical hedge portfolio at least to start with. Whether you actually hedge or not is a different issue, and that's really one for individual management to decide. But, clearly, if you decide not to hedge, it would seem to make more sense that you should be holding more in the way of additional provisions and more in the way of capital.

To summarize, I'd say that option pricing has a number of advantages when applied to these contracts. It minimizes or eliminates the effect of parameter sensitivity and parameter uncertainty. It gives you values that respond in an economically sensible fashion in response to changes in fund unit prices.

If you implement an option-pricing approach using a binomial or trinomial tree, then you have the further advantages of computational efficiency, provided you're only looking at a single stochastic risk factor. It also gives you a very effective and powerful way of modeling the reset feature. The major qualification is that these models establish the fair value of the contract and that, in practice, you would need to supplement it with additional provisions and additional capital.

CHART 1
WORLD INTEREST RATES 1/31/99

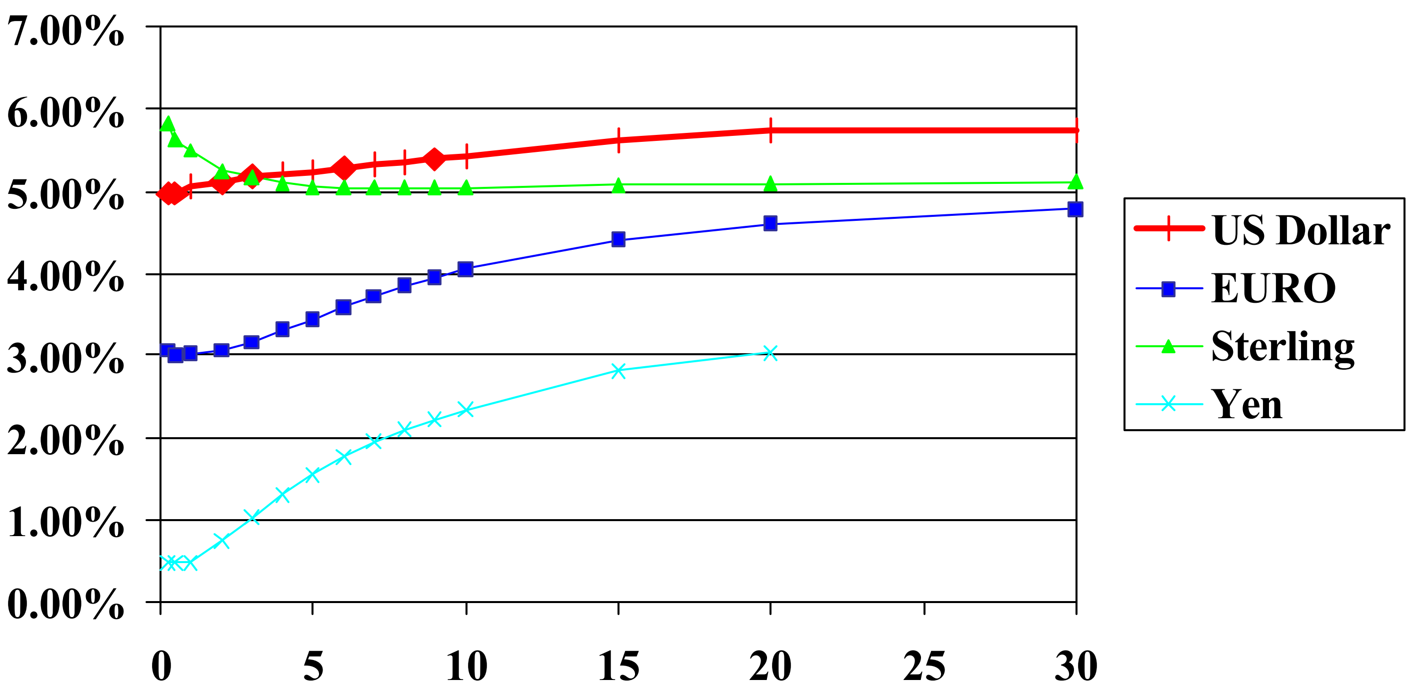

CHART 2
YIELD CURVES WITH SHORT RATE = 5%

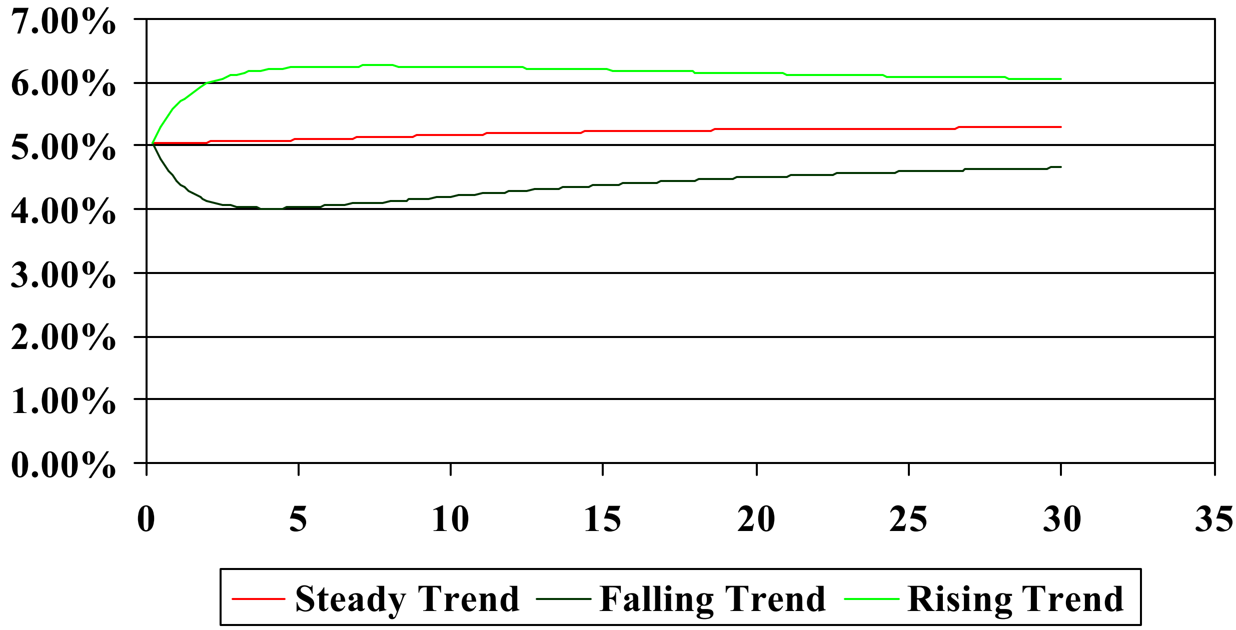

CHART 3

Value of Annuity Purchase Guarantee (Credit Guaranteed Rate)

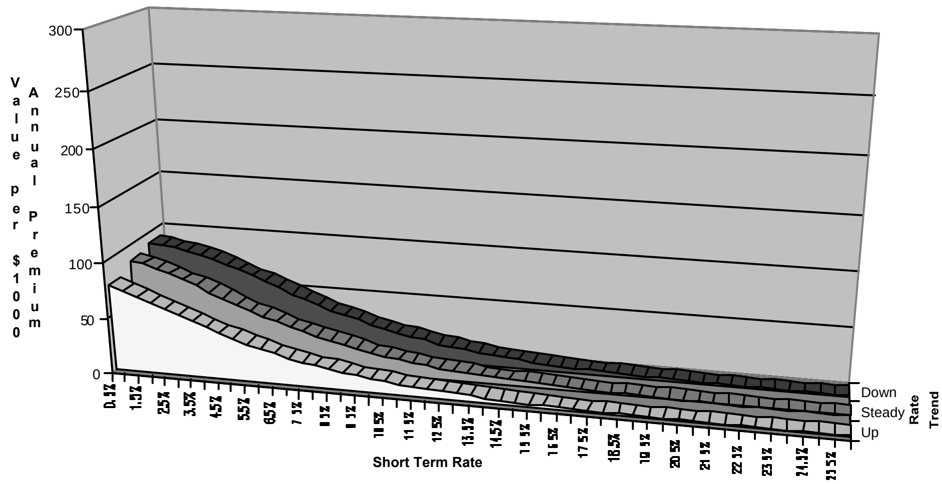

CHART 4

Value of Annuity Purchase Guarantee (Credit New Money Rate)

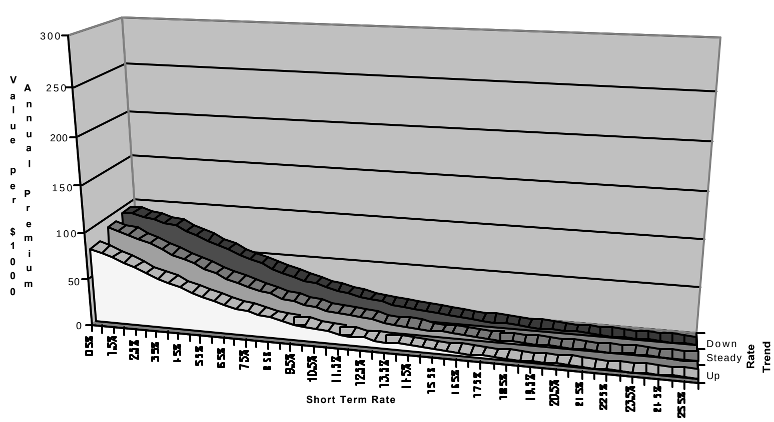

CHART 5

Value of Annuity Purchase Guarantee (100% Premium Persistency)

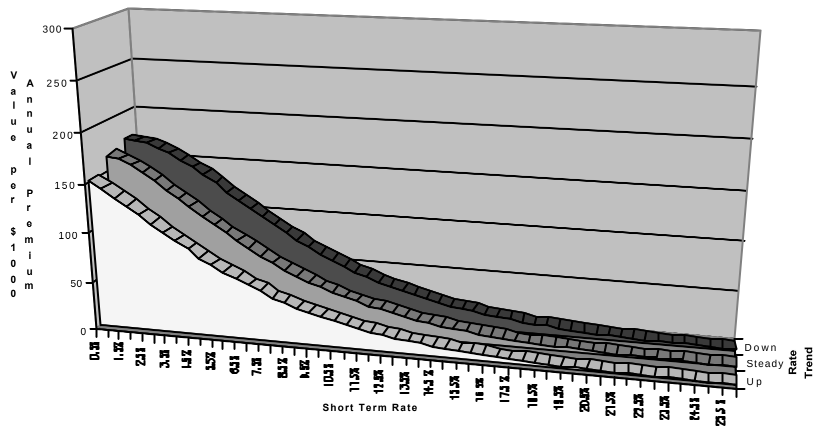

CHART 6

Value of Annuity Purchase Guarantee (Antiselective Premium Persistency)

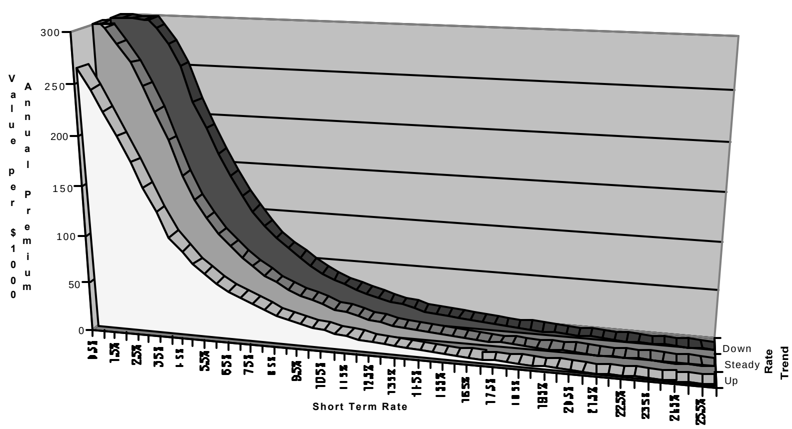

CHART 7

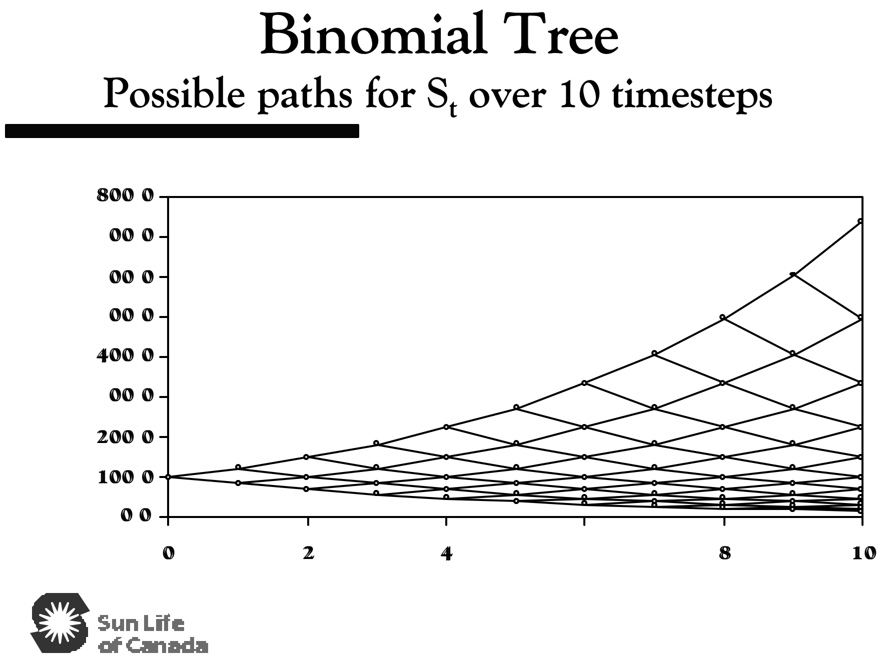

CHART 8

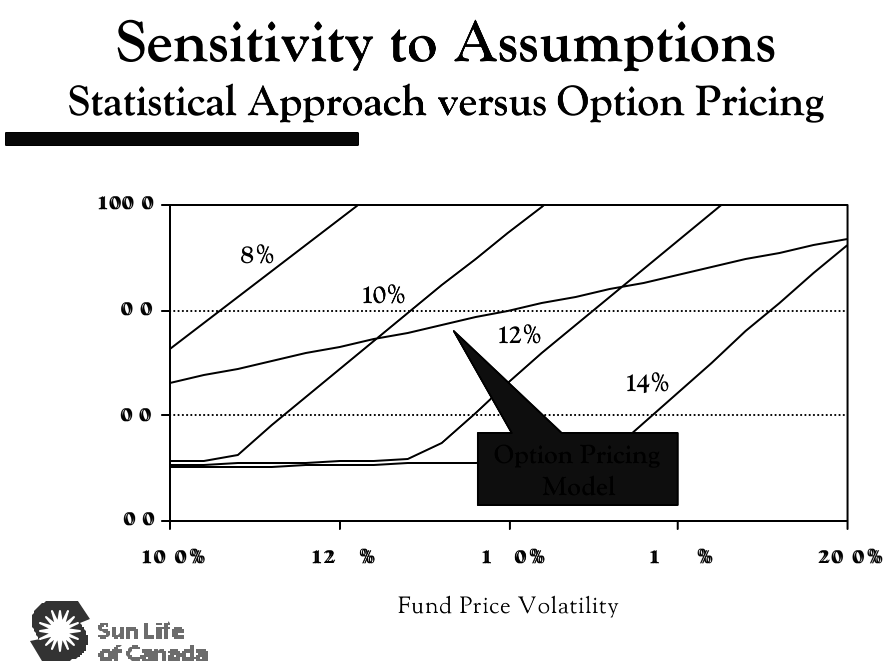

Text box reads "Option Pricing Model"

CHART 9

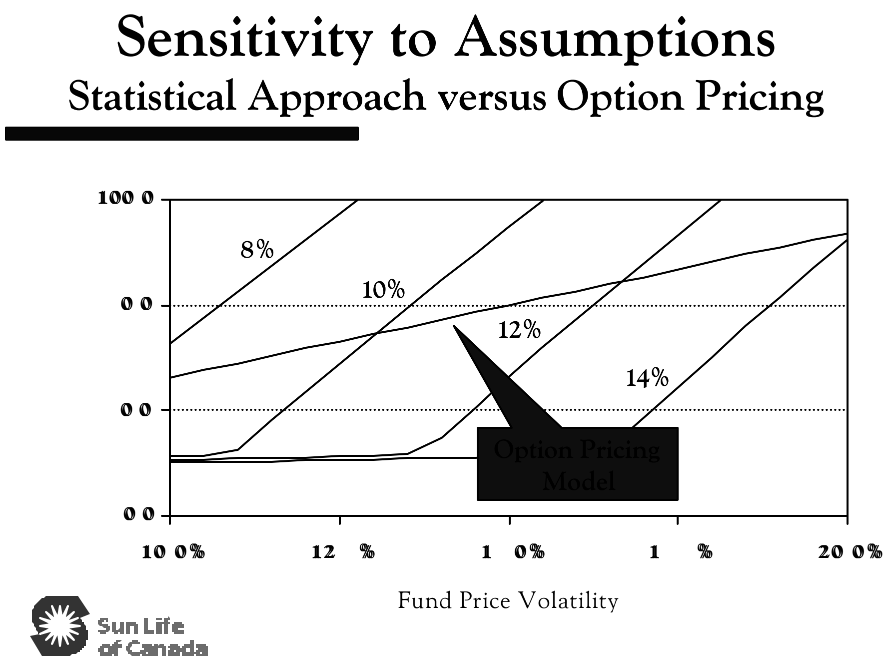

Text box reads "Option pricing model"