RECORD, Volume 27, No. 2*

Toronto Spring Meeting
June 20–22, 2001

Session 29PD
Valuing Guaranteed Minimum Death and Maturity Benefits: U.S. and Canadian Perspectives

Track: Canadian Institute of Actuaries/Financial Reporting

Moderator: SIMON R. CURTIS
Panelists: GEOFFREY HENRY HANCOCK
DANIEL P. MACKENZIE
CHARLES DANA TATRO

Summary: There has been an emergence of a wide array of minimum-investment performance guarantees for variable/segregated fund products in Canada and the U.S. In Canada, the Canadian Institute of Actuaries (CIA) has responded by developing a detailed framework for using stochastic modeling to establish the appropriate valuation liability for Canadian GAAP/statutory reporting. In the U.S., a number of techniques have been used for statutory and U.S. GAAP valuation.

This session describes the stochastic modeling framework developed by Canada for valuation and the practical issues encountered in its implementation. An update on developments in the valuation of similar features in the U.S. is also provided to allow a comparison of the approaches being taken in the two countries.

MR. SIMON R. CURTIS: We have three speakers who can give you a very good overview of the topic. I'll give you a brief introduction, and then we'll get right into their presentations.

Speaking first will be Geoff Hancock, a Fellow of the Canadian Institute of Actuaries (FCIA). He was very involved in the task force in Canada that developed a research paper on using captive techniques to price and value guaranteed minimum death benefits (GMDB) and guaranteed minimum maturity benefits (GMMB) in the

Canadian context. The paper has now worked its way into the mainstream of Canadian actuarial practice and actuarial liabilities in Canadian GAAP.

Second, we'll have Dana Tatro speaking. Dana is in a unique position. He works for a multi-national company based here in Canada, but he works in its U.S. operations. Dana has seen how to apply the Canadian standards to U.S. products because we have to do consolidated worldwide reporting on our U.S. business. Dana can give you an interesting perspective on Canadian and U.S. benefits and how the Canadian methods may or may not be applicable.

Finally, Dan MacKenzie is the appointed actuary at one of the largest writers of GMDBs and GMMBs in Canada. He will give you an overview of some of the implementation issues that have been faced by actuaries in Canada in trying to use stochastic methods to value these benefits.

MR. GEOFFREY HENRY HANCOCK: I'll be discussing valuing GMDBs and GMMBs and offering the Canadian perspective. I'll give a bit of an overview of this section by going over some background and scope for the benefit of our U.S. colleagues. I will talk a bit about the regulatory environment in Canada and how that led to the development of recommendations by the CIA task force. I will also talk about some of the highlights from that report, which is available on the CIA Web site. Then I'll talk about policy liabilities and get right into the valuation issues. I'll talk about principles, standards of practice, and some considerations that we discussed in our paper. I'll be placing particular emphasis on margins and provisions for adverse deviations—something near and dear to our hearts. I'll speak briefly on some other valuation issues including risk management and approximations in the valuation. Finally, I'll conclude with a few brief remarks on balance-sheet integration. I will touch on where we're headed in terms of capital and what the current requirement is.

First I will provide some background. "Segregated funds" is the common term used for these products, but of course they are really individual variable insurance contracts (IVIC). The U.S. equivalent is a variable annuity with guaranteed benefits. These are separate account products with the assets supporting the liability at market—the host contract liability. In Canada, we have some prescription on the guaranteed benefits that are offered on these products, so they can be classified as insurance and not securities. The guaranteed benefit must be 75 percent on return of principal on death and maturity, and the term to maturity cannot be less than 10 years.

These have been around for 30-plus years, but only in the last six years or so have these really gained in popularity—so much so that the growth in this market has out-stripped the mutual-fund industry for the last five years. The year 1995 was pivotal, particularly when the guaranteed-investment fund designs hit the market. This is where we saw guarantees being wrapped around brand-name mutual funds, the elective reset option coming into play, and the guaranteed-benefit levels being

at 100 percent return of principal, not 75 percent. So the guarantees became much more generous.

What happened? The industry as a whole became concerned about this. In 1997, the CIA formed a working group of which the formation was prompted by the fact that there were no industry guidelines or standards of practice in respect to the valuation. There were no prescribed minimum-capital requirements. There was a wide variety of practice in the market from doing nothing to doing some modeling. In retrospect, I think we would describe that as unsophisticated in terms of what we're doing today. Pricing was very much market-based and not model-based, or classified as very much market-based. You could classify it as inappropriate or inadequate in light of the information we have or the techniques we're using today, and there was very little risk management.

There was a joint working group that led to a research paper, which was a very good introductory document to the broad issues of risk in the financial management of these products. It really articulated the issues and then went on to recommend some direction for the CIA and the regulators in terms of developing standards. At about the same time, the CIA Committee on Life Insurance Financial Reporting (CLIFR), discussed valuation issues in its 1997 fall letter to appointed actuaries in Canada, but didn't really offer any further guidance other than to highlight the issue. However, in 1998, some real guidance came out, and there was a prescribed test that was offered—prescribed scenarios that were of the drop/recovery-type tests—single-scenario tests that would vary by type of fund. These were implemented to set a minimum-floor liability. Stochastic testing was still recommended to test the adequacy of these liabilities, because they were only meant to be minimums.

There were some key conclusions to the work that was being done at this time. The first was a recognition that deterministic methods—whether they were retrospective accumulation of margins or prospective single-scenario drop/recovery-type scenarios—were fundamentally flawed, that they were stop-gap measures only, and that stochastic modeling was the way to go. Then, in the fall of 1999, the CIA formed a task force entitled the Task Force on Segregated Fund Investment Guarantees.

The scope of our paper, which came out last August, was to develop and recommend approaches for using stochastic methods and stochastic techniques to measure the obligations. Originally, our focus was strictly on the valuation—setting policy liabilities and negotiating the implementation of those standards. But it quickly became apparent to us that, for these types of low-frequency, high-severity-type risks presented by the guarantees on these products, most of the provision actually would be in capital and not in the policy liabilities. It naturally led us to consider an integrated approach, so we looked at total-balance-sheet provision.

There were some working principles behind our rather lengthy report, which included:

1. A firm belief that stochastic techniques were superior to other techniques.
2. Consistency in practice.
3. Ease of implementation of any recommendations we might make.
4. An integrated approach to setting the provisions on the balance sheet.

For those of you who have not read the report, the key recommendations are:

1. The practitioners—the actual appointed actuaries having to set the policy liabilities—should be permitted by year-end 2000 to use these techniques and best practices as outlined in the report.
2. These models can introduce complexity in the valuation process. One of the things we realized within our mandate was to try to introduce some guidance that would lead to a narrowing of scope and range of practice rather than a widening. We attempted to do that for calibration and so forth.
3. We came out strongly and took the position that we should immediately discontinue the use of deterministic scenarios to set policy liabilities. We still think they have a very good place in capital resiliency testing, dynamic financial analysis (DFA) or dynamic capital adequacy testing (DCAT).
4. For year-end 2000, regulators should bring in a standard that would set minimum capital requirements. This is the factor-based approach that many of you are aware of.

The report is a comprehensive document. It addresses many things, including the valuation, investment models, modeling considerations such as policy-holder behavior, and so forth. It also addresses incorporating policy features.

The general approach to valuation that we take in Canada is the Canadian Asset/Liability Method (CALM), which is a purely prospective approach that's cash-flow based. The objective is to determine the basket of assets that are necessary and sufficient to mature the obligations. This can often lead to an iterative process in which you're trying to solve for that basket of assets.

The assumptions are "best estimate," based on emerging and expected future experience. However, because it is a policy valuation, there are provisions in the liabilities for adverse deviation (PADs). They can be accommodated in a couple of ways:

- To reflect uncertainty in the underlying assumptions.
- Through margins by modifying the underlying assumption or determining a provision in aggregate as part of the valuation process. The emphasis here is on GAAP accounting.

- The margins or provisions are not meant to be solvency based. They are meant to cover plausible adverse experience, not catastrophic experience.

Tied into all of this is the responsibility of the appointed actuary to set the valuation assumptions, ensure the methods are appropriate, and so forth. He or she must opine on the appropriateness of the valuation and the liabilities, and there are standards of practice and guidance notes to aid in that.

How do we value these guaranteed benefits? There is a contract out there—the assets of which are at market—so it's essentially a mutual fund with a guaranteed benefit attached. We take a bifurcated approach to valuation, which means it's not "whole contract," and that's what we've recommended in our paper. We are concerned with valuing the obligation created and the risk associated with the guaranteed benefits on their own. We're not modeling the whole contract. We use stochastic methods to model those guaranteed benefits, which are essentially embedded-put options.

One of the questions that comes up in our standards of practice and in the valuation in general is the term of the obligation. What is the term? The term of the host contract is typically zero. The assets are carried at market, and the liability is at the same value. For the guaranteed benefits, however, you need to extend the term beyond zero in order to measure the obligation. The key decision is what to do at the first maturity date of the product if the policyholder has the option to renew. That has a key impact on the results.

The last point relates to our integrated approach. The valuation is integrated with required capital. By that I mean we are living in a regime where we're determining the total-balance-sheet provision using factors. The actuary would determine the policy liabilities. The difference is minimum-required capital. Setting of margins for adverse deviation or provisions for adverse deviation should not take that into account. It should be related to uncertainty. They're not supposed to be related to a company wanting to manipulate capital and reserves.

Here are the seven steps that a company would need to go through to value these benefits:

1. Select a stochastic model for the investment returns.
2. Calibrate the model according to criteria that we've proposed in our report.
3. Generate scenarios.
4. Incorporate margins for adverse deviation into the assumptions that are not scenario tested.
5. Start with current policy values, current guaranteed amounts, and current market values for the underlying funds, and then value the benefits along each scenario consistent with that scenario.
6. Determine the liability for that scenario. Typically, this means discount the

guaranteed benefit cash flow to the valuation date.

7. At the very end, establish provision for investment risk. This last step is one of the primary decisions faced by the appointed actuaries.

I'm not going to talk about models per se, but in our report we briefly address them. We talk about calibration—the key point being we do not exclude. We felt it was not appropriate or necessary to either exclude or endorse any models, so that's how we proceeded. However, we did realize that we wanted to introduce some type of narrowing range of practice or a sufficiently narrow range, so we have these calibration requirements. Since we're focused on the left tail of distribution, when the markets go down, there are adverse consequences when there's a payout under the maturity or death benefits. We focus on the left tail, and that's what our calibration is looking at. It's ensuring that your model produces sufficiently adverse events with enough frequency.

How did we develop this calibration table? It is a table that gives you some numbers. Where did it come from? It came from modeling. There were three people on the task force that contributed to a fair amount of modeling in coming up with those points. It was based on the TSE300 Total Return index for the last 45 years, and there are three different models that led to the same relative conclusion. We didn't just add them, average them, or pick the lowest or the highest. The results were subjective, and we looked at the results between the three models and came up with a table. When an actuary is going to implement his or her model, it can be any model, as long as it satisfies this calibration when it is calibrated to or fitted to the TSE300 over that same time horizon. If it does not, the actuary must make adjustments to his or her model parameters so that it does pass the calibration by increasing the volatility, decreasing the drift, or by doing both. If you have a more complicated model, it can be done by adjusting other parameters. We give some technical guidance to this in the report.

There's no specific guidance, however, on projecting fixed-income returns. We're talking about equity returns in terms of calibration. It's important to note that this calibration process itself is just to ensure that your model is producing fat-enough tails. It's not meant to be a margin.

You would apply margins to non-scenario tested assumptions such as risk factors that are presumed to be known in the valuation or can be reasonably expected to be scenario/path-independent. Mortality is a good example. There are certain experience risk factors that we know are probably going to be path-dependent, but we model them as if they were path-independent. An example may be lapse rates. You need to be careful in setting the margin to these non-scenario-tested assumptions. The standard range is five to 20 percent, plus or minus.

The sign of the margin is particularly important should you decrease lapse rates or increase lapse rates to increase the liability, because that is the purpose of the margin. It's to adjust the base assumption that results in an increase in liability. It's

not always immediately clear or obvious at the valuation date whether or not that margin—the adjustment to the base assumption—should be positive or negative. It can change by scenario. It can change by duration. So, some sensitivity testing might be needed.

Lastly, in terms of non-scenario-tested assumptions, the most important test usually ends up being surrenders and withdrawals—the rate of fund depletion. It's important for the actuary to test this assumption and understand the interaction between the guaranteed benefits that are paid out and the revenue stream that underlies the fund.

These are the assumptions that are not scenario tested and that vary by scenario. Normally, this means some sort of dynamic behavior, and I'll give one approach here that you could take to that. I think this is one that works well within the guidance and within existing standards, and that's to break your assumption to a static and variable component. A static component is non-economically related. It's an underlying assumption that would apply in a stable or average market over the long term. It would include a margin. The variable component would allow the end result—the resulting assumption—to float above or below that static level. It would be based on what's happening in the scenario, and it could have a performance component. These are options we're trying to value. The policyholder holds an option. It can have an option value or a time component, market-value component, and a performance component. The key suggestion in our report, and what's actually being done in practice, is that the appointed actuary need not, and indeed should not, assume the extremes of behavior—neither complete irrationality or 100 percent efficiency. Underlying all of this, you do need to consider and have due regard for reasonable expectations.

One of the major decisions facing the actuary is how to accommodate a provision for the investment risk—the market risk. This is done through a measure called the conditional tail expectation (CTE). You run your stochastic model through 1,000 scenarios and get 1,000 results. Which one do you choose? We deviated from percentages, something that we're all comfortable with, and felt that the CTE measure was more appropriate for a couple of reasons. At level X, CTE is nothing more than the weighted average of your valuation results beyond the X percentile. You order your results from smallest to largest, and take the average of the tail. Why did we feel that was a good measure? It's robust, which means it's less sensitive to statistical fluctuation in terms of the number of scenarios where you have a limited amount of time, so you're testing a limited number of scenarios and small changes in the market value to guaranteed-value relationship at times zero. The same cannot be said for percentiles or mean-plus-deviation measures. It always reflects down-side risk, so whatever CTE level you choose—whether it's CTE (70 percent) or CTE (80 percent)—you're always taking the average beyond that point and it always considers the entire tail. That's important because it's consistent with the methods that were used to develop the factors underlying the total balance sheet provision, so that's a plus. Although it's slightly more difficult to calculate

than percentiles, it's still easy to calculate and it's readily interpretable.

One point I did not mention is that this valuation is done under a realistic set of scenarios using the realistic or the actuarial approach to valuation. It is not a risk-neutral valuation.

What is this investment /market risk? There are three components to it that we described in my report:

1. Investment performance.
2. Parameter misestimation—parameters in your stochastic model.
3. Model risk.

The accepted range that we recommended in our report and that has come to be a standard right now is CTE (55 percent) to CTE (80 percent). More practically, it's CTE (60 percent) to CTE (80 percent) I haven't seen too many companies sitting at less than CTE (80 percent) Remember that CTE (70 percent) is approximately the 85^{th} percentile.

Now I will discuss investment model uncertainty. Even if you knew the model that governs stock returns and you knew the regime and parameters—whatever they might be—there's still uncertainty because there's a random component in the model. That's the whole point of using these. You need to establish some absolute lower bound to account for that uncertainty. We chose CTE 50 percent. Take the average of the worst half, not the better half. We started from there.

The next thing you need to recognize is that even if you know the model, you don't know the model parameters. So, even if someone gives you the perfect stochastic model and it does describe how the markets behave, you don't know the parameters for it into the future. We generally use historic data to estimate the parameters, and there's statistical uncertainty in those parameters and standard errors associated with those parameters, which are sometimes very large. You need to reflect that in the provision. It is more pronounced for longer-term contracts that are out of the money. For contracts that are in the money, there's less uncertainty associated with the model parameters because you're more certain of the payout.

There are some more sophisticated techniques. We're actually examining this type of uncertainty. They're called Markov-Chain Monte Carlo (MCMC) techniques. They're a Bayesian approach, which essentially recognizes that the parameters are random variables. Hence, you assume a distribution for those random variables. Then you provide some simulation to simulate those model parameters because they're random variables. You input those into your stochastic model and generate scenarios. Although it's a very effective process, it's difficult to implement and time consuming, but it is much better and much more rigorous than ad-hoc sensitivity testing.

The last risk that we want to incorporate in this provision—and remember, the provision is determined by where we set the CTE level at the end of the day—is model risk. Your model is not perfect. You know that the log normal model doesn't accurately describe how the market has moved. No one really knows how the market has moved. You might not be accommodating everything in your model. You might want to have a monthly model, but you don't. You have an annual model. So, there are things of that nature. It's trying to check the correctness or the stability of the model through time. It's also meant to account for the fact that you have a limited amount of time, so you can't run an unlimited amount of scenarios. Also, of course, there's uncertainty in client/agent behavior such as resets, lapses, and transfers.

Another thing our report addresses is that, on the one hand, you're trying to project the out flow—the cost of the guaranteed benefits. Under CALM, it is a present value. It's an out-flow less/in-flow approach. The revenue here is not the total-management fee charged from the fund. There are claims to that revenue stream, so you have to attribute a portion of that revenue stream to the guaranteed benefits. There are claims to that revenue stream such as investment management, other expenses, and so forth. It's not a trivial task to actually take the total-revenue stream and attribute it to its various components. Remember, we're taking a bifurcated approach to valuing this guaranteed benefit—not a whole contract approach.

This is one approach to doing that. The key point though, is that the piece of the total-management fee that you attribute to funding the guaranteed benefits does not necessarily have to relate to pricing loads. It's not that obvious and not necessarily that clear. Another key point is that whatever amount you attribute to funding the guaranteed benefits and that you're going to account in the valuation, you have to forego in the host-contract valuation. You can't double count.

Risk management should be reflected in the valuation. Reinsurance, hedging, and so forth—to the extent that they're actually in place—are actually implemented. You can't reflect risk management or risk mitigation for practices that aren't in place. Of course, there are various degrees of risk management such as doing nothing, running the risk naked, buying bonds, full/partial hedges, static or dynamic, and of course, reinsurance. There are additional risks and uncertainties imposed by these practices and you need to establish PADs for them. Straight pass through or risk transfer to reinsurance might introduce some slight credit risks and liquidity risks. It's debatable. For other strategies such as hedging, there's basis risk, of course liquidity risk, and rebalancing uncertainties in the future in terms of volatility and correlations in the market. You need to reflect these on the valuation. PAD

One point is that risk management in and of itself might not lead to a lower reserve. It actually could lead to a higher reserve, so you need to keep that in mind. It isn't necessarily a bad thing. It could lead to a more stable reserve through time, and hence, less earnings volatility. But it's unfair to jump to a conclusion that risk

management in and of itself will lead to a lowering of reserves. One of the key things our regulator, the office of the Superintendent of Financial Institutions (OSFI), has said is that in order to give credit to risk-mitigation practices such as hedging, they want to see a well-defined risk management culture within the company. It can't be something that you're implementing on an ad-hoc basis. It has to be part of the culture—well documented and well audited.
The first of the approximations is discounting. The rate we use to discount the cash flows is fixed. It would not vary by scenario. It needs to be appropriately conservative and include a margin. These stochastic models are complicated to implement in their own right. But of course, when you're dealing with a huge portfolio of 250,000 contracts, the in force population data takes on very special meaning and it has a special place in the valuation.

Although preferable, sometimes seriatim valuation isn't possible. If you're going to take approximations such as grouping contracts together, you need to use special care. You can mask or distort the risk exposure by grouping contracts that are not similar. Another piece that comes in are the funds. Many companies offer dozens of funds, which are not obviously linked or tracking an index. Of necessity, the valuation will need to come up with proxy funds that are meant to mimic their own underlying segregated funds. There is risk there—basis risk and model risk.

Now I will address a couple of brief valuation issues. There are a number of scenarios you're going to use along with the time step. They have a big impact on the results. There are trade-offs between the accuracy you get and the resources that you have. Transparency of the valuation process is important for auditing purposes. The scenarios should be reproducible, auditable, and so forth. But the volume of input—numbers such as the in force data—can make this problematic.

I can leave you with one point where we stand—add complexity only where it's warranted. I think most companies will agree with this philosophy. It may at times seem that we're going overboard in bringing in these models, but I would highly advise that bringing in too many moving parts may obscure the analysis. Keep it simple, and gradually make it more complicated when warranted. Focus on the things that matter such as your stochastic model. It doesn't need to be perfect. Get a good model and look at persistency. Don't neglect one aspect of the risk valuation for another. Total liability in Canada here must be non-negative in order to have your negatives entered for some contracts if you're doing this on a seriatim basis. But at the end, when you set your CTE level at the confidence level you deem appropriate, if that number is negative, you would bring it up to zero.

The results can lead to significant earnings volatility, which is no surprise. It's a very practical implementation issue. Absent any other changes such as small changes in market value and/or keeping the CTE level the same, could lead to large increases in the reserve. Earnings volatility is really on the front burner.

I'm not going to describe the entire factor-based approach that we have to total-

balance-sheet integration, but it is the factor approach. It's a reasonably complicated but straightforward set of tables that we apply to come up with a total-balance-sheet requirement. The actuary would then deduct the policy liabilities and hold the remaining piece as minimum-required capital. That's where we stand. The regulator has indicated that it would like to move toward an environment where companies would be permitted to use models to establish liabilities, as well as capital, by using an integrated approach. We're all looking forward to that day.

A practical issue to keep in mind is the expense of keeping a dollar provision. This is an integrated approach—a dollar provision in capital versus a dollar provision in liabilities. It's much more expensive to keep it in capital. It's a very important decision for the appointed actuary and a great responsibility to determine that the liability that's set up on the balance sheet is appropriate, given the heavy pressures that he or she will have on managing capital and the volatility of earnings. There are some real practical issues here.

Lastly, we will discuss DCAT. We still think it has a very useful place in looking at these products. What if testing still has a place, and the actuary that has a significant amount of business on the company's books is looking at single scenario such as drop-in recovery-type scenarios in trying to understand the company's exposure, when it comes time to do the annual DCAT report?

MR. CHARLES DANA TATRO: I have the great honor of being a U.S. product actuary that gets to do both Canadian and U.S. reserves. It is fun. I don't know how much of an honor it really is, though. We're going to go through a little bit of what I get to go through on a daily basis. First of all, I'll talk about the differences between the U.S. and Canadian benefit designs. I think my discussion will be beneficial when we actually get to the reserving methodologies in order to see why the regulators chose the methods they did. Then we'll go into a brief comparison of commissioner's annuity reserve valuation method (CARVM) reserves compared to the latest OSFI guidelines that came out. Then we'll get into some numbers and briefly touch upon a few developments in the U.S. in regard to separate account (SA) guarantees.

First, we'll talk about death benefits. Death benefits in both the U.S. and Canada are very similar. You have the basic return of premiums (ROP), periodic ratchets, annual ratchets, roll-up-indexing death benefits, and the bull and bear, which is the greater of the roll-up or the indexing death benefits. There is not very much of a difference there.

We see a bigger difference when we get to guaranteed minimum income benefits (GMIB). In the U.S., income benefits have been gaining popularity since 1998 when the Equitable first came out with theirs. Now you can get income benefits that provide ROP, periodic ratchets, annual ratchets, roll-ups or indexing, or the bull and bear, which is the greater of the roll or the annual ratchet. In Canada, the income benefit never really took off. They lean more toward the guaranteed minimum

accumulation benefit (GMAB), where you provide a 75 percent to 100 percent guarantee after a 10-year waiting period, which may or may not have a discretionary reset. Those accumulation benefits never really took off in the U.S. A few companies offer them, but they usually have restrictions on the funds that you can go into—restrictions to exercise them, and it didn't really seem to be a benefit that a lot of clients were looking for when they were buying a variable annuity. The U.S. method for base reserves is CARVM. I'm sure most of you are familiar with that. It's a prospective formula-driven approach, which is very conservative. You can't use lapses. You have to use provisions for adverse deviations (PADs) assumptions. There's no company-specific deferred acquisition cost (DAC). It's all driven by the level of your surrender charges. They like you to do it on a seriatim basis. It gives you the present value of your highest benefit whether it be a death benefit, income benefit, or annihilation surrender, and you get the whole reserve. CLIFR is the prospective cash flow approach. You use realistic assumptions with PADs. You do have a company-specific DAC, and you can use some portfolio-basis assumptions to make your valuation a lot quicker.

The GMDB reserve is where we really see the differences coming out. In the U.S., it's a deterministic approach. You basically have five asset classes: a drop, a recovery, death benefit costs, present value, back to the future—and then you have your reserve. The difference between your base reserve and your GMDB reserve is held in the general account, and there's no integration with your reserve and your RBC, which is similar to the way it is in Canada with your reserve and your capital requirements. Geoff did a pretty good job explaining the Canadian methods, so I won't go through them in as much detail as he did.

The U.S. method for the GMIB and the account value benefit reserves are now leaning toward the stochastic approach. That's being proposed by the Variable-Annuity Guaranteed Living Benefit (VAGLB) task force. I'm sure a lot of you are aware of that. If you have a very simple benefit that you can analyze, you can use either a stochastic approach or the Keel method, which is just a simplified stochastic approach. Again, it uses the same five asset classes that they use for the death benefit, however, this time you have to do it stochastically. They're not going to tell you which drops and recoveries to use. You actually have to get a log-normal model, they'll give you the input parameters, you project out 1,000 scenarios, and then you do the present value of your reserves. Again, the difference in the base reserve and this reserve is your reserve is held in general account and it's not integrated with RBC. In Canada, they use the same method for their death benefits that they use on their income and their account value benefits, which is mainly the stochastic approach using the TSE300 as the base to validate whether or not the model is correct. It can then run any asset classes once you have a validated model. Your reserve is the present value of your benefit, less any margins that you have available to cover that benefit. Again, your reserve is set anywhere between CTE 55 percent and CTE 80 percent, and then the minimum continuing capital and surplus requirements (MCCSR) is the difference between your CTE 95 percent and whatever your reserve happens to be.

Now I'll discuss some reserves that actually show numbers. I ran a pretty simplistic model—the S&P 500 only—a stochastic basis using the Canadian approach and also by using the drop recoveries for the equity class. I did a male age 60 at issue and used MER to 250 basis points. For U.S. actuaries, an MER is your total-asset fees including fund-management fees, M&Es, and any distribution-related expenses that you have. I assume no margin for offset here, which is a pretty big assumption, and it's going to cause a big difference between the U.S. and the Canadian reserves. I'll talk about that when we see the numbers, because that margin for offset can significantly reduce your requirement on a Canadian basis. Both the GMIB and GMAB that I did have a 10-year wait. The GMIB suggests an annual ratchet, so you increase it to the maximum anniversary value. The GMAB is a 100 percent guarantee after the 10-year waiting period.

Table 1
Comparison of Reserves

U.S. CARVM Reserves

		Death Benefits			GMIB	GMAB
Date	**S&P**	**ROP**	**Annual Ratchet**	**5% Index**	**Annual Ratchet**	**ROP**
12/31/1995		1.12	1.12	1.12	13.97	0.00
12/31/1996	23.1%	0.00	1.21	0.13	54.93	0.00
12/31/1997	33.4%	0.00	1.31	0.00	64.64	0.00
12/31/1998	28.6%	0.00	1.40	0.00	72.35	0.00
12/31/1999	21.0%	0.00	1.38	0.00	79.68	0.00
12/31/2000	-9.1%	0.00	3.31	0.00	78.90	0.00

Candian OSFI Reserves

12/31/1995		4.16	22.35	19.26	23.94	8.34
12/31/1996	23.1%	1.41	29.03	13.36	31.93	4.35
12/31/1997	33.4%	0.32	40.69	6.29	46.20	1.91
12/31/1998	28.6%	0.12	54.73	3.13	63.82	0.70
12/31/1999	21.0%	0.07	69.11	2.03	81.64	0.08
12/31/2000	-9.10%	0.12	70.77	5.20	82.78	0.13

Table 2

U.S. CARVM Reserves

		Death Benefits			GMIB	GMAB
	S&P	ROP	Annual Ratchet	5% Index	Annual Ratchet	ROP
12/31/1969		1.12	1.12	1.12	13.97	0.00
12/31/1970	4.0%	1.09	1.21	1.51	54.93	0.00
12/31/1971	14.3%	0.22	1.31	1.07	64.64	0.00
12/31/1972	19.0%	0.00	1.43	0.22	72.35	0.00
12/31/1973	-14.7%	0.62	3.83	2.84	79.68	0.00
12/31/1974	-26.5%	5.15	19.03	32.22	112.47	85.58

Candian OSFI Reserves

12/31/1969		4.16	22.35	19.26	23.94	8.34
12/31/1970	4.0%	4.13	24.54	24.78	26.99	10.58
12/31/1971	14.3%	2.22	29.48	23.04	33.47	8.97
12/31/1972	19.0%	0.98	36.69	18.05	42.79	5.73
12/31/1973	-14.7%	3.32	38.28	46.46	43.17	19.80
12/31/1974	-26.45%	20.44	57.84	113.54	83.29	124.51

I found this comparison (Table 1) to be pretty interesting. The first column is the date and the second column is what the S&P 500 returns were for that time period. The different benefits go across the top. One thing I found very interesting on the CARVM approach was that your ROP death benefits and your annual ratchet death benefits start out at the same reserve. Since you're using a deterministic drop in recovery, your annual ratchet always looks like an ROP. The CARVM doesn't seem to capture the little nuances of your benefit, because it's always a straight drop and then a recovery. The five percent index on a CARVM basis looks just like your ROP on an annual ratchet for the first year, and then you start getting into different reserves going out in the future, which are basically influenced by what your actual fund performance was. On your GMIB, which is an annual ratchet again, you see significant reserves being put up on the VAGLB guidelines. That's mainly because of the conservative assumptions that are used. You can't assume any annuitization or lapses, and all your annuitizations occur after your waiting period if the benefit is in the money. Having 100 percent utilization of the benefit and no lapses is a pretty conservative reserve holding for a GMIB. Your GMAB comes out pretty good because of the time period I chose. I have a 23 percent return for five years. You're not going to have a GMAB that's in the money under any reserve basis if it's just an ROP.

On the Canadian approach, because it's a stochastic analysis of the actual cash flows and the actual benefit, you see that each of the death benefits puts out a different reserve. This is because the stochastic investment scenarios that you use capture the little nuances of your benefit—whether it's an annual step, a pure return of premium, or a five percent indexing in there constantly driving up your death benefit value.

The first thing that jumped out at me when I looked at this was the $22 per thousand per annual ratchet under the OSFI guidelines, and $1.12 under the CARVM guidelines (Table 2). I had to think about that for a while just to make sure that it was correct. The reason for this huge discrepancy is that I did not assume any margins for offset on the Canadian basis. If you were to assume that you had 25 basis points of your total MERs to cover off this benefit, and say that the average length of your contract is eight years, then if you have a seven-year contract, and you have PADs that reduce your lapses and other aspects of it, that 25 basis points will be worth roughly two percent of your premium up front. That would be $20. If you reduce the $22 by $20, you're down to a $2 reserve margin for offset, which is pretty powerful. As Geoff alluded, if you use that to reduce your reserves in capital under the OSFI guidelines, you can't use it in your other valuations. Five percent indexing gets captured very well using the OSFI guidelines. The same is true with GMIB, because you're able to use much more realistic assumptions, and your GMAB gets captured a little better. It causes you to put out a reserve, even though in these scenarios you wouldn't have to put up a reserve under the stochastic analysis. It shows that you should be holding something in case you have that one-time margin shot that affects your benefits.

I also did this analysis using a different S&P 500 period with a few bad years, mainly 1969–1974, in order to test whether or not each of the reserves reacted once the account value dropped. They did react, and this will point to the earnings impact that you can have on one of these benefits. The year 1973 has a 14 or 15 percent market drop. Your reserves increased significantly under both bases, which will make it very difficult to manage your earnings on either a U.S. or Canadian basis. The GMIB also increases under both scenarios. This is where your GMAB really causes you some pain, especially under the CARVM approach. This shows you don't throw off any benefit until the fifth year, where all of a sudden you've got to throw off $84 per $1,000 of space that you have, because you had two bad years of performance that wiped out all of your existing PAD that you had because of the market. Your reserves started jumping up well before you actually hit that period because you had that scenario where things got progressively worse. You captured that in your reserves.

Now I'll cover a couple of other developments in the U.S. There is FAS133, and I'm sure any of you out there with reinsurance transactions on GMAB are having pain with this. It basically says that you have to do the fair market value of any embedded derivative in your contract. Death benefits are excluded because it has a

life-contingency piece. GMIBs are usually excluded, unless you reinsure it and the reinsurer pays on a lump-sum basis, not on a periodic basis. Then the reinsurance transaction actually creates an embedded derivative that you have to value on the FAS133.

I'm sure we're going to see RBC requirements out there for SA guarantees. We don't have them yet, but they've got to be coming. We've also seen some new GMDB designs in the U.S., mainly an earnings-enhancement benefit that was first developed by Sun Life. This benefit is very interesting. It doesn't move the same way that traditional benefits move. If the account value is down, this benefit doesn't pay out because it pays out on the amount of earnings you have in the contract in order to offset taxes at death. When you're having good years, and account values are positive, if you have a client that passes away, this benefit would pay out where your other benefits may not, because the other benefits are meant to protect against a down-market scenario. This one is to give the client a boost when he or she has an up-market scenario.

This new design isn't handled very well by either the current Canadian or U.S reserve methods. It is definitely not handled well by the U.S. reserve, because as I said, it's all drop scenarios. If you're always dropping your account value, this type of benefit isn't going to be in the money and you're not going to put up a reserve under CARVM. Under the Canadian methodology, you would put up a reserve. However, you usually get out-weighed by any reserves that you have to put up for your GMIB or any of your more traditional death benefits—mainly because this is a newer benefit, and you don't have the same amount in force to be an overriding factor.

MR. DANIEL P. MACKENZIE: I'm going to discuss some of the implementation issues with respect to stochastic valuation. The topics I want to cover are mainly with respect to Canadian GAAP, but there is a lot of overlap for FAS133, which is an emerging topic in this as well. I'm going to touch on data verification, fund classification, model calibration, scenario generation, policyholder behavior, setting recoverable margins, and margins for adverse deviations.

With regard to data verification, data integrity is critical. This product is really transactionally intensive, and there's a lot of backdating and key-plus-three and transfers. One thing my firm relies on is getting two monthly extracts—a "snapshot" and the transaction records. We try to reconcile the two. We build what we call a transactional account value and compare that to the "snapshot" account value. More often than not they agree, but every month there are a couple of policies in which they've done something special to it that has to be investigated. To us, this is a first step. We don't go past this step until we get a sign-off on what we call the transactional account value.

I prefer a theory seriatim valuation. There have to be some compromises in this process in terms of grouping. I prefer to group the assumptions, because in the

end, if I have a theory seriatim valuation, it provides drill-down detail. It also gives insights into the relationships between your assumptions of policyholder behavior and the liability.

Below is a checklist of what you would need for basic data requirements:

- Snapshot:
  - age, gender, market value, guarantee value, investment fund, product code, distribution channel, maturity date(s)
- Transactions:
  - deposits, surrenders, transfers, resets

Our actual extracts would have more producer information and other interesting things to look at. But in terms of just valuing the liabilities, you need the items in this list.

With regard to fund classification, Geoff did mention that there are a lot of funds including managed funds, index funds, and balance funds. They all have their own investment objectives and historical returns. It is an interesting process to try and project the investment returns for what is really a proprietary fund. There's no other fund like it in the world. To give you an idea of what I've been up against, our company has gone from having one fund in 1994 to having 17 funds two to three years ago, and now we have over 60 funds. I hope the growth is not so exponential in the future. The method that we use to assign or project investment returns includes taking each fund and assigning a proxy benchmark, but typically it's a weighted average of proxies. We talk to the investment manager. We look at the historical returns and try to get a nice fit between the proxies, which are basically external-market entities. Then each proxy is assigned model parameters (e.g., mean and volatility) based on the historical returns.

After classifying your funds, the next step would be calibration. Essentially, we take the task force approach. The approach outlines the so-called tail calibrations where you have some strength of the various percentiles. We replicate that process or the various proxies that we have, and monthly returns are used to provide a larger sample size. In general, a monthly projection of investment returns will also be more indicative, more robust, or less sensitive, if you will.

Correlation of returns by asset class is an interesting issue for implementation. It's quite difficult to come up with a correlation matrix, and it is more relevant to the actuarial liabilities than the capital. In the capital scenarios, basically all asset classes go down the drain, so they're very correlated in the tail. To me, the correlation should be dynamic and not static. For example, you could treat historical correlation as being expected, but there's diminishing probability of uncorrelated adverse returns, so the worst things are the worst things all around. It would be conservative to assume returns are uncorrelated.

Investment model scenario generation is where you're trying to balance computer run-time with reliability of results. In other words, you want to be able to look at

something instead of your computer running all night. My preference is to use a large number of scenarios applied against a small, but representative sample of contracts, so you can run all of your scenarios, but not necessarily against your entire in-force. Then you would rank those scenarios by the liability and select a few representative scenarios that you would then apply against the entire in-force. This approach would require sensitivity testing to ensure that a choice of "large," "small," and "few" would be appropriate. It is actually representative of a full valuation or a more rigorous, detailed valuation. The other approach you can take is grouping the data into cells and running all the scenarios against your group data. If you're going to be doing that, you want to group by age, gender, term-to-maturity, investment class, and the market value/guaranteed value (MV/GV) ratio.

The projection period should be near the contract maturity date. Most of the segregated funds would have what are called deposit maturity dates, which would occur before the end of the contract. In essence, it's really one of the embedded options that the policyholder has the option to renew at the end of the maturity date. So you want to be valuing that option as well.

Now I will discuss policyholder behavior and how to reflect that in the valuation. In regard to lapses and actual experience, the existing data, which is about five-years-old now, may be sufficient for establishing a base lapse rate for the early years of a contract. But for later years, your lapse experience is pretty scant, and you really should be relying on guidance. There is some guidance that essentially says you should be conservative. One thing to note if you do have a lot of registered retirement income fund/life income fund (RRIF/LIF) business, which is a tax-driven payout annuity contract in Canada, is that there will be some regular lapses due to the payout.

•Lapses
In terms of reflecting lapses in your scenarios, the lapses should be dynamic and in sync with the scenario-specific investment returns. In-the-money contracts should be less likely to lapse and even less likelier when they're nearing their maturity guarantee date.

•Sensitivity Testing
The results will be quite sensitive to lapses. It's probably the most sensitive policyholder behavior assumption of them all, and this really underscores the need for sensitivity testing of your lapse assumptions.

•Resets and Actual Experience
There is a lot of actual experience on resets, but the current data probably reflects the growing level of client and broker sophistication. It is likely that reset utilization will rise in the future. But when you're looking at historical reset behavior, you have to look at the opportunity to reset. That has to be considered, because in the last year or so, there's actually been limited opportunity to reset. It would be erroneous to say that resets are down. They're non-existent for an underlying reason. This

speaks to Geoff's point about 100 percent utilization and 100 percent sophistication. Most clients with elective and/or limited resets do not reset at the peak. They can't time it perfectly. They can see it's going down, and they reset on the way down. Or they reset on the way up, and now they're out of their contractual number of resets.

In terms of rolling this into a scenario, a practical approach might be to assign a probability of reset, so it's going to be zero when the guaranteed value is less than the market value. That's the time factor. It will be zero when there is no opportunity to reset. You would have a maximum when the market value is just slightly in excess of the guaranteed value. It's going to be low again when the market value is well in excess of the guaranteed value. Essentially, these are people who have decided to hold on. They don't want to reset because that will push out their maturity date, or they just don't know what they've bought. There are contracts with limited resets, so it is important to reflect the underlying contract options and the limitations. For example, there are contracts where you can only reset twice a year, so you wouldn't be resetting this client six times a year in your valuation. Consistent with your reset assumptions, you've got to reflect the contract, and that's a tough thing to model. You have to get that data when they reset, then reflect that and capture it.

•Transfers and Switches

When people switch from one fund to another, you have to look at the contract as well. There are contracts with some inhibiting mechanisms for transfers. It triggers a reset and things like that. This has to be considered as well. There really are small proportions of clients that are active traders right now that would switch in and out of their favorite funds. None of them can effectively time the market. They're day traders and they're not going to beat the market or the company. They're not going to be able to select against you. Their impact, in terms of what they can do to you by adding value to their options, is negligible. There are a few traders that are using dollar-cost averaging. They're putting their initial deposit into a money market fund and then drawing it down every month.

If you wanted to put some of these assumptions into specific scenarios, the current experience indicates that the transfers would increase with market volatility. These are people who are trying to time the market. The current level of activity does not necessitate that transfers be reflected in the valuation. There's just not enough of it to be significant. One thing you might want to consider, in terms of scenario-specific assumptions, is where selected contracts are at or are near in relation to the money when they're approaching maturity. The person might be tempted to switch to a riskier, potentially higher-returning fund, and they essentially have a one-year guarantee at that point.

•Additional Deposits

Additional deposits are usually optional at the client's discretion. What you see in terms of patterns of additional deposits can be categorized into three groups:

1. There are a lot of single one-time deposits. These are people who found money or inherited money, didn't want to put it in the bank, and put it in a seg fund.
2. There are a lot of people who are depositing an amount annually, but it's differing amounts—lump sums. Basically, it's what they can afford.
3. There is another group that would have monthly regular deposits, as per some retirement plan that they've established for themselves.

In terms of putting this into scenario-specific assumptions, my tendency is not to model it, because additional deposits really lower the guarantee liability. They're adding to their principal, and most adverse scenarios would only help you in terms of reducing your liability. Some contracts allow additional deposits in the years just prior to maturity. So essentially, instead of having a 10-year guarantee, they have a more valuable one- or two-year guarantee. They may exercise that option to give you additional deposits, and you might want to do some sensitivity testing to see if that's material for you.

•Miscellaneous Policyholder Behaviors

Mortality—the results are not particularly sensitive to the mortality assumption. As Geoff touched upon, mortality is one of the things that is not directly correlated to the economic scenarios. But in terms of picking levels of mortality, general population mortality is probably more suitable than typical annuitive mortality. There may be some correlation between resets and mortality. People in poor health may be more likely to reset, so you might want to consider some kind of mortality spike after a reset.

Another thing to look at in terms of policyholder behavior is contract size. Larger contracts typically are more sophisticated in terms of behavior. They're going to reset and lapse more effectively.

I'll move now to the setting of recoverable margins. I will provide some discussion on the breakdown of fees. Expense fees are used to offset fund-related expenses, such as audits and regulatory filings. Management fees are used to pay for a number of things including:

- Maintenance and investments management expenses. Typically, these would be expressed as a flat percentage or basis points of assets.
- Acquisition expenses. Within the valuation, you must demonstrate that that portion used to amortize deferred-acquisition costs is going to be sufficient to do so.
- Guarantee costs. You may want to take special care with respect to the portion used to cover income taxes. Give some consideration to whether or not that's going to be available to cover the guarantee costs, because in most scenarios, you're not making any money, so you're not paying any taxes. That might be one way of looking at it. Getting at this breakdown is not straightforward.

Pricing assumptions could be suitable depending on how "fresh" they are and to the extent that all aspects of product costs have been delineated. That's been my approach.

Margins for adverse deviations—I tend to split this into two categories, policyholder behavior and the investment model. The margins for policyholder behavior assumptions should really be based on sensitivity testing in terms of seeing which ones are really material in the valuation and which ones you have to give a lot of care and attention to. As a practical matter, you might allow the sign of the MfAD to be constant, since it probably will be for most of the adverse scenarios. Making the sign more dynamic is a tricky matter and requires a lot more programming than it will give you in terms of insight and value. The lapses should produce the largest PAD and then be followed by resets.

In terms of the investment model or the margin for investment returns, the MfAD should be incorporated into the choice of the CTE level.

Earnings volatility has a close cousin now called captive volatility, and there's no mechanism within the valuation approach to handle earnings volatility. My opinion is that there should be some smoothing for the maturity, which is a long-term benefit. You definitely have volatility with respect to mortality benefits, but it's just something that companies are wrestling with. They're wrestling with it in both earnings and capital volatility at the same time. These could be viewed as symptoms of a hedging program, where if you're hedging the economic risks, you're also going to be hedging some of this earnings-and-capital volatility.

Another item that is possible is the reserves versus capital issue. If you're holding capital at a multiple, and you're not holding actuarial liabilities at a multiple, the question is, which is more conservative—to have a higher reserve or not? Because to the extent that you have a lower reserve, you're going to have more total-balance-sheet requirements. It's an interesting question.

Finally, I'll provide a summary of the critical issues:

- The data verification is crucial. It's the underpinning of both your assumptions and your results.
- Calibration is the key to a good investment-return model.
- Run time needs to be balanced with the reliability of results.
- Lapses and resets are the most significant among policyholder-behavior assumptions.
- Avoid double counting when establishing your recoverable margins.

MR. DAVID FIHRER: (Canada Life) Geoff, you mentioned using the reinsurance to offset the reserve. Getting more to what we're discussing here, and Dan alluded to this as well, the paper uses an eight percent fixed lapse rate. We saw how significant the effect of dynamic lapses is. Doesn't it make sense to take that into account? It seems possibly even more significant than the CTE level.

MR. HANCOCK: Yes, lapses have a significant impact on the final results. Ideally they should be dynamic, but you need to take care in establishing the relationship you're going to build into the valuation methodology. Assuming a 100 percent efficiency in policyholder behavior, I firmly believe it is unrealistic and shouldn't be done. Nor should it be based on complete irrationality. It's problematic postulating a realistic relationship. I agree in theory that it's a good thing to do. In practice though, it can be difficult. Is it the largest component of your provision? Perhaps—but perhaps not. You need to take great care in establishing an appropriate sign for your lapse margins or in making sure the provision is positive, because there's an interplay between the cost of the guaranteed benefits, which in and of themselves are lapse supported. Lower lapse rates lead to higher payout on the guaranteed benefits. But of course you have the recoverable margins, the margin offset, or the attributed revenue stream to fund it. It's not always clear that a lower lapse rate leads to a higher policy liability. You need to take that into account in your dynamic relationship. It's not trivial. I think this is one of the main reasons preventing everyone from jumping to a dynamic lapse level. We're not recommending using fixed rates, or using dynamic.

MR. FIHRER: Implicit to the fact that we have to now use the fixed lapse rates, you're not saying we should be testing for sensitivity?

MR. HANCOCK: In setting the policy liabilities, yes.

MR. CURTIS: I tend to agree with that. I'm wearing my financial-reporting-chair hat for the CIA. I would think that you would if your models are capable of doing it. You would want to test lapses dynamically. I think when we had the task force earlier, to a certain extent you have to walk before you can run. I think it was recognized that indiscriminately throwing dynamic behavior into the models would introduce a lot of volatility into the results, where people were not at the stage of being able to understand it. Certainly in our own company, when we put in dynamic assumptions, they do have a material impact on the reserves and generally will obviously lead to higher reserves if you're trying to put in anti-selection. I think it's a matter of moving in any kind of steps that you can handle, because the stochastic approach has the potential to overwhelm the person running the models. You get more numbers than you ever dreamed of, and quite often you have trouble interpreting results, even with static assumptions.

MR. FIHRER: I have a few more comments regarding the available offsets and the significant effects of that. We've had some discussions, and the paper doesn't give much guidance on how to actually come up with this offset on the capital side and the reserve side. There could be more guidance. We've seen a lot of companies doing significantly different things. We don't want to be in a position to do something and someone else comes along and tells us to do something else and we're in a rather awkward position. We like it to be discussed at a committee level.

The last thing I've got is regarding DAC. The issue now is the effect of the term of

the liability being zero and whether or not it's going to produce a negative reserve. All of this is great in order to find out what your underlying reserves are, but you're going to be reporting a zero reserve, which effectively means if you can recover your debt, then you're holding a reserve in essence, which is bigger than your policy premium method (PPM) reserve and your guarantee reserve. There are a number of issues here. The first issue is reporting a best estimate, then a PAD, and then an extra reserve. How are we going to report this? Do you have a scope in order to say that you've changed your assumptions, however, you've got a higher guarantee reserve, which means a lower capital level? Could we do that? That's playing games, but that's the essence. What does that actually mean for this business? Running these stochastic models for no guarantee reserve, and then going back to the fact that we've got no reserves doesn't make sense. Does that mean you've got no offset on your capital? Again, we feel there's not much guidance. It is a little bit subject to interpretation.

MR. CURTIS: I'm not going to answer the question in detail, because there are a lot of questions. But I'll point out that that initial task force report from last summer was really on how to build, calibrate, and develop a model in order to model the benefits. It was recognized, but it left a lot of unanswered questions in terms of when you went to practically implement these. You just gave a laundry list of practically every issue that we've come up with. The Financial Reporting Committee will be issuing a research report the summer of 2002 that will actually address every one of those issues. In some cases it may still be a bit vague on what the right action is, but they've all been identified and there will be discussion on all of those issues and the current thinking of how they should be addressed. Things such as bifurcation versus whole-contract valuation, the fee-offset issue, how you manage volatility, whether or not you can build up cushions, strong performance, integrating it with DAC recoverability testing, and the appropriate PAD ranges will be covered. At this stage, it's going to be more similar to actuarial literature, but our hope is that we can format it as education notes, so it starts to refine some safe harbor practices.

FROM THE FLOOR: I have a question on stochastic scenario generation. Why use historical results, instead of a more implied volatility capital market approach?

MR. HANCOCK: We discussed that in our task force. The main difficulty is that using typically implemented risk-neutral valuation makes some assumptions about creating a replicating portfolio hedge. Now, the underlying assets aren't tradable, and creating a hedge is problematic. When using a straight-risk neutral calculation such as taking a log-normal model and integrating it with Black-Scholes, just use the risk-free rate of return, run a log normal model, and take the expected value. Leaving aside what the volatility number is, that's risk-neutral,. It's a big problem, because that assumes a replicating portfolio. Of course, you're creating a head on an untraded asset. How do you account for that, and how do you account for transaction costs? How do you set reserves and capital using a stochastic model when you get one number? I don't know. Do you just change the volatility and run

it twice? I don't know. We took the approach to be more consistent with Canadian practice, which was to use a realistic model for all your assumptions, best estimate, prospective, cash-flow based, and apply margins.

FROM THE FLOOR: Is that realistic if you are trying to test the left tail?

MR. HANCOCK: I think it is. I always look at risk-neutral as well. I've seen the risk-neutral results. It usually comes in around the 90th percentile—maybe a little higher for an unhedged liability. I think that using a CTE (80 percent) measure is pretty realistic. In terms of our total-balance-sheet requirement, I think that is higher than a risk-neutral result, and I think that's appropriate for an unhedged position.

MR. MICHAEL DUBOIS: (Mass Mutual) I was just going to one of the comments that Geoff had made regarding the methodology described as consistent with Canadian GAAP. I don't believe that U.S. GAAP has come up with any kind of standard method for handling this. I think they're still in the discussion phases. When you say consistent, is this the method that a Canadian GAAP would use? What types of differences do we see? Dana, do you have any comment on what types of things are being done for U.S. GAAP?

MR. HANCOCK: Yes, it is Canadian GAAP, and you have to keep in mind that we only have one set of financial statements, so it is stat and GAAP. When I said consistent, I meant it's consistent with the principles underlying GAAP. But for purposes of your company preparing its financial statements, this is the method used.

MR. TATRO: On a U.S. GAAP basis, there is no prescribed method yet, and actually companies are all over the board with what they do. A lot of companies just use a retrospective-premium accumulation. Some companies are using a stochastic analysis of the benefit, but I'm not really sure where they got that stochastic analysis or what kind of generators they're using. The methodologies are all over the board, and none of them have been ascribed or endorsed by FAS 133, so I'm not sure that a task force has been put together for the U.S. GAAP basis.

MR. MACKENZIE: Regarding the previous comment, I know that on FAS 133, it's really the market value of those embedded options and market values are not defined as Black Scholes or risk-neutral, so it's a vague definition. In my opinion, this is an approach to market value, but I think that's an open question.

MR. MARK KINZER: (Clarica) My incomplete understanding of this whole process is that if you have a hedge, you give some sort of credit against the reserves and/or capital. I'm wondering if you'd like to comment on it, and whether or not there are any changes foreseen in this summer report that's coming out?

MR. CURTIS: I'll comment on the reserves. Right now in the reserving piece—the

actuarial liability—you can take credit for hedging. Again, it's not an area that we have spent a lot of time with regarding the list of issues we're going to address, because it hasn't been identified as an issue for reserving. It is an issue for capital, because right now the capital formula is factor-based, and as such, it's a very unfriendly vehicle for trying to develop any type of credit for hedging.

MR. MACKENZIE: When they formally introduced the factors in November 2001, they did commit to providing the rules for getting capital offsets. I understand those are in draft form right now, and there's heavy emphasis in the risk-management culture. There's a lot of discipline and documentation that needs to be in place before you get the credit. I assume it will be an evolving process. You wouldn't get what would be your ideal credit offset—certainly not for this year-end or even next year. It will move along over the next three years. That's the time frame.

MR. FIHRER: Some of the models you've looked at have an annual generation and some have monthly stochastic generations. For example, Mary Hardy's model does monthly, but doesn't do interest rates, it only does equity yields. The Wilkie model does both, but it only projects annual rates. Have you done any testing to see if there's a significant difference between annual versus monthly? We've got a significant balanced fund, so we know when you have to generate the interest.

MR. MACKENZIE: My preference would be a monthly model. Given the way the tail is being calibrated, I think you probably want a multi-parametric model that has the best fit possible.

MR. HANCOCK: I agree with Dan. I have investigated the impact of the cash-flow time step because, of course, periodicity for your investment return model can differ from your cash-flow model, as long as they're integrated. I think what you really mean is your cash-flow model. You could be sampling monthly returns from your investment model but then, for pragmatic reasons, using an annual time step in your cash-flow model and actually valuing the liabilities. The time step is important for some contracts. It's really that simple. It's really how subject you are to intra-year fluctuations. Products that have resets are more sensitive.