RECORD, Volume 31, No. 1*

New Orleans Life Spring Meeting
May 22–24, 2005

Session 34 Seminar

Stochastic Embedded Value and Its Use in Risk Measurement and Financial Management: Part 2

Tracks: Risk Management, Financial Reporting

Moderator: HUBERT B. MUELLER

Panelists: MICHAEL A. HUGHES
PENNY COULTHARD†
MARIA TORRES-JORDA

Summary: Recently, there has been increasing interest in the subject of integrated financial and risk management in the United States and in Europe. This session explores the use of stochastic embedded value (EV) in financial management and risk measurement. Topics addressed include the considerations in reconciling period-to-period results as well as expanded case study examples. Each of the sessions rely on the use of simple but relevant case studies to illustrate the stochastic EV mechanics, its use in risk management and the reconciliation process. The risk measurement discussion includes the constraints to the risk management process.

MR. HUBERT B. MUELLER: Many of you have come back from the first session, and maybe some of you are here for the first time this afternoon. This session is to provide more in-depth case studies on the topic of stochastic EV, both from a real-world and from a market-consistent perspective. For those of you who were not here before, we discussed some of the topics in the prior time slot.

My name is Hubert Mueller, and I'll be your moderator. In this session, we're going to have three speakers. The first two will be from Ernst & Young, and then the third

† Ms, Penny Coulthard, not a member of the sponsoring organizations, is senior consultant at Tillinghast in London, United Kingdom.

speaker will be from Tillinghast's London office. We'll start with Maria Torres and Mike Hughes from Ernst & Young.

Maria works in Ernst & Young's insurance and actuarial advisory services practice in New York. She has experience as a chief actuary at a life insurance company in Argentina, and she is now working as a consultant in various areas, including statutory and U.S. GAAP financial reporting, profitability measurement, reinsurance, valuation, economic and EV frameworks and fair value. She is a frequent speaker at industry seminars on economic measurement frameworks and the actuarial aspects of the international financial reporting standards (IFRS). She's also a member of the Society's IFRS task force.

Mike is a senior actuarial advisor in the Chicago office of Ernst & Young. He has about 18 years of industry experience. He's a recognized expert in the area of financial reporting and performance analysis and does have extensive experience in the area, including mergers and acquisitions, financial modeling, actuarial process improvement, asset/liability modeling (ALM) and litigation support.

MS. MARIA TORRES-JORDA: In this presentation we are going to illustrate, by means of a case study, the different methods considered by the CFO Forum for valuing options and guarantees. We are first going to briefly introduce each of the methods conceptually, and then we are going to get into the case study, which is based on a variable annuity (VA) product. At the end of the presentation, we are going to extend the lessons learned from the case study to other products, and we are going to discuss practical issues and steps that need to be considered when implementing stochastic EV frameworks.

The CFO Forum considered three approaches to value options and guarantees. One approach is stochastic on a real-world basis. In this method, the traditional form of EV incorporates a time value of financial options and guarantees by taking the expected value of the distributable earnings over a range of stochastic scenarios. The scenarios are based on real world outcomes for economic assumptions such as interest rates and equity returns.

The second method is the certainty equivalent, or risk-neutral, method. It is also known as market-consistent embedded value (MCEV). In this method, investment returns are projected at the risk-free rate, and the discounting is also done at the risk-free rate. The value of the options and guarantees is performed using a risk-neutral valuation.

The third approach that the CFO Forum considered is a combination of the traditional deterministic EV, in which cash flows are projected using the expected future return on each asset class and discounted at the risk discount rate, with the time value of financial options and guarantees calculated separately using a risk-neutral valuation.

After analyzing these three approaches, the CFO Forum finally decided to go with the first approach, i.e. stochastic on a real-world basis. As I briefly explained before, the EV reflects the expected value of future distributable earnings across a range of realistic stochastic realistic scenarios for interest rates and equity returns. Some companies choose to discount the different scenarios using the same risk discount rate. Some of the companies vary the risk discount rate by scenario, setting the discount rate equal to the risk-free rate for each scenario plus an equity risk premium.

The stochastic on a real world framework has its advantages and disadvantages. Let's talk about some of the advantages first. It is capital asset pricing model (CAPM) based, which is consistent with how companies manage their business, by setting target returns higher or equal to the return on capital required by shareholders. It is consistent with the way that companies price their business and perform actuarial appraisals. It incorporates complex long-term interactions between assets and liabilities, shareholder cash flows, and management actions. It also reflects regulatory constraints, which can take the form of conservative reserving or minimum capital requirements.

However, the framework does have several disadvantages. One of its key disadvantages is that it can result in mispricing of assets and EV anomalies. By projecting the investment returns based on the expected asset yields, and then discounting back at the risk discount rate, traditional EV would tend to place higher value on business backed, for example, by corporate rather than government bonds. Companies need to be careful when setting the risk discount rate to appropriately reflect the risk of the assets backing the business.

Another disadvantage is that the framework is not used by financial markets to price risk. They tend to use risk-neutral methods. Also, the risk discount rate is very difficult to fine tune for business units / products and asset portfolios. Companies tend to use a top down approach to estimate their discount rates. The starting point is to determine the shareholders' required return using CAPM. If the company is partially debt financed, the required return on equity is then combined with the company's cost of debt to compute a weighted average cost of capital (WACC). Additional adjustments should be made to come up with the required return for each business unit and product lines. But since, based on market data, we can only compute quite objectively the beta for the entire company, coming up with a risk discount rate by product usually involves significant subjectivity and assumptions. In addition, as explained before, it also is difficult to calibrate the risk discount rate so that a change in investment strategy does not have an impact on the EV net worth.

Let's move now to the certainty-equivalent, or risk-neutral approach, which is basically what we have been referring to before as "market-consistent EV." This framework addresses one of the key criticisms to the traditional EV methodology. The traditional EV framework can result in mispricing of asset and liability cash

flows by just using one single risk discount rate, but this framework addresses this issue. Let's explain it with a simple example.

If, for example, you decide to invest $100 in equities and the return on equities is 8 percent, at the end of the year you would expect to have $108. If, alternatively, you decide to invest in bonds, and you put in $100 with a 4 percent expected return, at the end of the year you will have $104. If the risk discount rate is not determined appropriately in your traditional EV calculation, you may end up placing a higher value on the money invested in equities, which would be incorrect. At time zero, both investments are worth $100.

The basic principle of the certainty-equivalent, or risk-neutral framework, determines that you have to discount each of the cash flows based on a discount rate that reflects the risk inherent in the cash flow. For example, in this case, you will discount the money that you invested in equities at 8 percent and you will discount the $100 that you invested in bonds at 4 percent. Each of the investments will be worth $100 at time zero.

But how do we do this in practice, and what do we mean with the "certainty-equivalent method"? In practice, to make things simpler, companies tend to risk-adjust the cash flows and discount all the cash flows using the risk-free rates. This is just a mathematical device that allows you to overcome the problems that you may face if you have to estimate the risk discount rates for each of the cash flows.

You may be thinking that risk adjusting the cash flows can be as difficult as estimating the correct risk discount rate. In practice, it is quite simple. If you are dealing with cash flows that do not have market risk (for example, fixed cash flows or cash flows that only have pure insurance risk), you don't need to do any kind of adjustment. For example, for a term product, you will directly discount the cash flows using the risk-free rate.

Cash flows that depend on asset returns are adjusted so that the underlying asset returns are equal to the risk-free rates. Projecting and discounting at the risk-free rate will give us the same answer as projecting and discounting at the actual asset rate. But projecting and discounting all assets (no matter how risky they are) at the risk-free rate, imply that we do not have to worry about the risk discount rate to use for each projected cash flow.

There's more to MCEV. No investor is going to assume the business just by getting in exchange the present value of future liability cash flows discounted at the risk-free rate. They will require an additional allowance to account for what is called the frictional capital costs, which include: double tax, agency costs and the company’s own credit risk.

Let’s discuss the frictional capital costs in more detail. What does "double taxation" mean? An insurance company resembles a leveraged investment fund, in which

debt is raised through the sale of insurance policies rather than through the capital markets. In contrast to investment funds, insurance company shareholders in most markets are liable to pay taxes twice on the investment return on assets backing required capital. Investment returns are first taxed when they flow through the insurer's taxable earnings and then again, as part of the shareholders' taxable income when distributed as dividends.

The other cost that also has to be accounted for is agency cost. This cost represents the additional return that shareholders would require to compensate them for the possibility that management may not always act strictly in their best interests.

Some people believe that you have to adjust the liabilities for the company's own credit standing or claims-paying ability. As with standard debt instruments, insurers have the option to default on their liabilities. Since policyholders may not receive the full payment entitled to them, this option represents an asset to the insurance company. It's usually reflected by increasing the discount rate. Instead of discounting with the risk-free rate, companies tend to discount with the risk-free rate plus the company's own credit risk.

This risk-neutral framework, like the stochastic real-world framework, also has advantages and disadvantages. Among the key advantages, we can mention that the framework is consistent with the financial market pricing of risk (i.e., it is consistent with how financial markets price tradable instruments). It minimizes the subjectivity in setting the discount rates, and it should properly value items that have observable market prices.

Among the key disadvantages is that it is not aligned (at least now) with how insurance business is being priced or managed. Considering that there's not a liquid market for the insurance liabilities, it is difficult to calibrate the value of the liability to market prices. Approximate valuation techniques need to be adopted to determine the frictional capital costs.

But as Hubert said, though this framework is very similar to the fair-value approach proposed at some point for IFRS Phase II, which companies opposed, many of the companies are now looking at this methodology, trying to implement it internally, and starting to manage their business along the lines of this framework.

Now I'm going to turn it over to Mike, who's going to get into the case study.

MR. MICHAEL A. HUGHES: Before we get going, let me quickly rattle through some of the details of the block of business that we're going to look at so that you can have a sense for some of the details of the case study. It is a VA block. We tried to make it reasonably realistic. It has five years of existing business. The account value on each of the years of existing business is about $100 million. There's also one year's worth of new business with $100 million of premium. All of the business is in the separate account; there's no general account fund. It's single premium

business, so there are no renewal premiums. This particular product has a guaranteed minimum death benefit (GMDB). It has a 5 percent roll-up feature, and the roll-up has a 200 percent of premium cap on it.

There are various product charges. There's a $30 annual administration fee, mortality and expense (M&E) of 150 basis points (bps) and an explicit guaranteed minimum death benefit fee of 25 bps. Surrender charges start at 7 percent and grade down gradually to zero. The management fee is 60 bps; that's really investment management expenses that are borne at the separate account level rather than an explicit charge. Commissions are 5.5 percent of premium. Trail commissions start in the second year at 1.2 percent and then fall down to 20 bps in renewal years. Acquisition expenses are $130 per policy and 2 percent of premium. Maintenance charges are $60 per policy and 24 bps on the account value. There are some transaction costs of $110 per death and $25 per surrender. I think the average size was $50,000.

Termination rates start low at 2 percent and grade up. They max out at 20 percent in the year following the surrender charge and then drop back down gradually to 10 percent. For partial withdrawals, we have 3.75 percent annual free partials. That might be a little high, but that's what we have here. Also, mortality is at 75 percent of the Annuity 2000 table.

For simplicity, statutory reserves are set equal to tax reserves are set equal to the cash surrender value. We have target surplus equal to 0.75 percent of the separate account assets plus 0.12 percent of the net amount at risk, and we have a 35 percent tax rate.

Let's start with a deterministic look at the world: a traditional EV on a deterministic basis. To do this calculation, we need to make two assumptions. We need the separate account return assumption, which we'll set at 8 percent, and we need a risk discount rate, which we'll also set at 8 percent. It's typical that you would have consistency between the rate at which you project the equity returns and the rate at which you would discount them back.

Next we'll look at the pattern of distributable earnings for one year's worth of new business. You can see in Chart 1 that there's a surplus strain in the first year, followed by gains thereafter. The gains are lower in the early years when the surrender charge is wearing off and cutting into profitability. The return on investment (ROI) is 19 percent and the present value of distributable earnings (PVDE) is $1.36 million. Because the return is significantly higher than the discount rate, you end up with a positive value for new business, which is what you would expect. But it's important to keep in mind that this is on a deterministic basis, so the death benefit guarantee is not really costing anything, making this somewhat an overstated view of profitability.

Chart 1
Distributable Earnings – 1 Year of New Business, Deterministic Basis

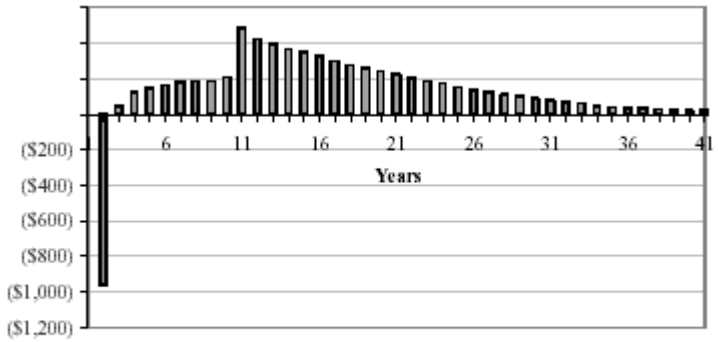

IRR = 19%

New Business EV = $1.36 million

Next we'll look at the EV reporting on a deterministic basis. Chart 2 shows the EV at year-end 2004 and year-end 2005. This is a balance sheet view in a traditional actuarial format. It breaks down the EV into the three different components: the target surplus, the present value of future profit (PVFP) calculated on a book-profit basis and the cost of capital. If you play with your algebra, the net of those three items is equal to the PVDE. In this case, there are capital transfers to or from the line of business so that there's no free surplus held within the VA line of business. You can see that in this case we have $14.2 million of EV at the end of 2004 that grows to $17.1 million at the end of 2005.

Chart 2

Deterministic EV Results – EV Balances

Year	2004	2005
Adjusted Statutory Book Value		
Target Surplus	$ 3,780,991	$ 4,344,061
Free Surplus	-	-
Adjustments to Statutory Capital and Surplus	-	-
Total Adjusted Statutory Capital and Surplus	3,780,991	4,344,061
Value of Existing Business		
Present Value of Future Profit	11,543,586	14,003,938
Cost of Capital	(1,127,290)	(1,254,137)
Total Value of Existing Business	10,416,296	12,749,801
Total Embedded Value	$ 14,197,287	$ 17,093,861

Chart 3 shows the movement in EV. Hubert touched on this in his presentation. There's an unwind of the discount on existing business. It's an 8 percent discount, so you get 8 percent return on the opening EV. New business contributed $1.36 million. That was measured at the point of sale, so there's a little bit of an unwind of the discount on new business. There's also a capital transfer in this example so that we end up with the right amount of required capital. There are no experience variations or assumption changes in this example. It just shows on a deterministic basis how we get from one period to the next.

Chart 3

Deterministic EV Results – EV Income and Movement (Millions)

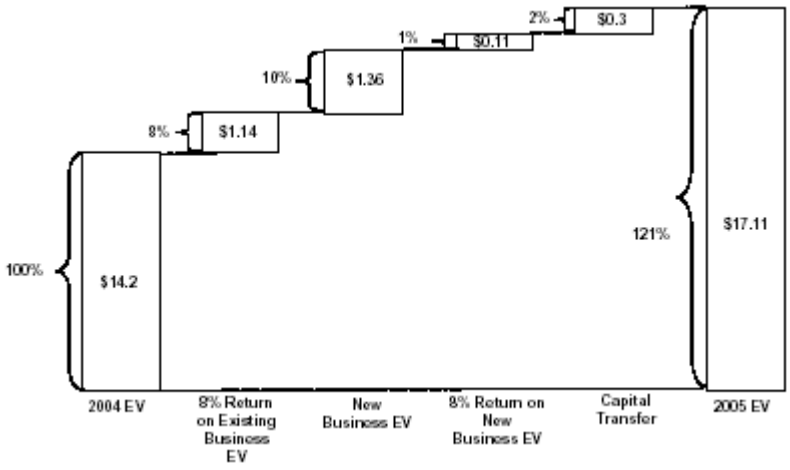

Chart 4 is a transition to the stochastic world. When we're doing our stochastic EV on a realistic basis, we're going to assume that the average equity return is 8 percent, and we're going to set the variability of equity returns using a lognormal model. The model has been parameterized to reflect historical equity returns. I think we have a 16 percent volatility assumption and an 8 percent expected return. This graph shows the dispersion of cumulative equity returns over time. I'm showing the 10th, the 25th, the 50th, the 75th and the 90th percentiles to give you an appreciation for how spread-out the equity returns get in the analysis.

Chart 4

Economic Scenarios – Real World

GMDB Benefits Capped at 200%

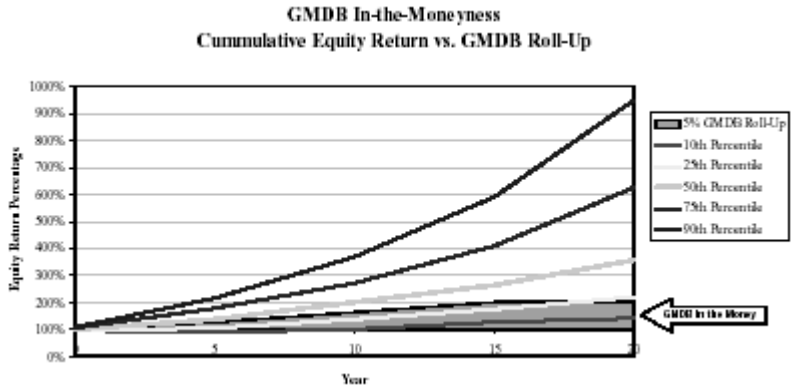

We've overlaid the 5 percent roll-up benefit onto those percentiles. That's the

shaded area at the bottom. You can see that the 5 percent roll-up maxes out at 200 percent and then stays level. I thought that this was an interesting way to give people an appreciation for what was happening in the stochastic scenarios and in what percentage of the scenarios the guarantee is in the money. You can see that the roll-up benefit is in the money in a little more than 25 percent of the scenarios.

Chart 5 shows a rank ordering of the PVDE. In the worst case, the PVDE is approximately a negative $20 million. In the best case, it's almost $40 million. The average PVDE in this case is $10 million. The deterministic EV was $14.2 million and the mean stochastic EV was $10 million, so the value of the options and guarantees, or the value of the GMDB in this case, was $4.2 million. That equates to about 15 bps annually as a percentage of the account value. So for European embedded value (EEV) reporting purposes, if you're going to report on a realistic stochastic basis, you'd report the $10.0 million as your EV, but you'd show it in the two components of the deterministic EV and the value of the options and guarantees.

Chart 5

Embedded Value under Real-World Stochastic Scenarios

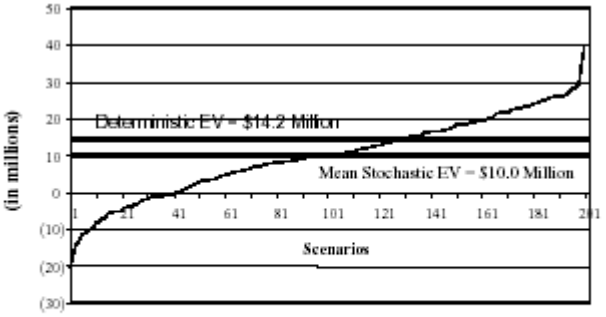

Deterministic EV	= $ 14,197,287	
Value of Option	= (4,159,283)	≈ 15 bps annually
Real World Stochastic EV	= $ 10,038,005	

It's also helpful to look at what's happening to new business. Chart 6 illustrates the distributable earnings on a stochastic basis as well as on a deterministic basis. These are the average distributable earnings by year. The strain in the first year is a little greater on a stochastic basis on average because in some cases you're paying those death benefits. The profits in the renewal years are also a little lower, again, because of the cost of the death benefit guarantee. In this case we used the average distributable earnings—this was a question that came up earlier—to come up with an ROI or an internal rate of return (IRR) of 14.1 percent. The value of new business now has dropped down to approximately $900,000.

Chart 6

Distributable Earnings – 1 Year of New Business, Stochastic Valuation, Realistic Basis

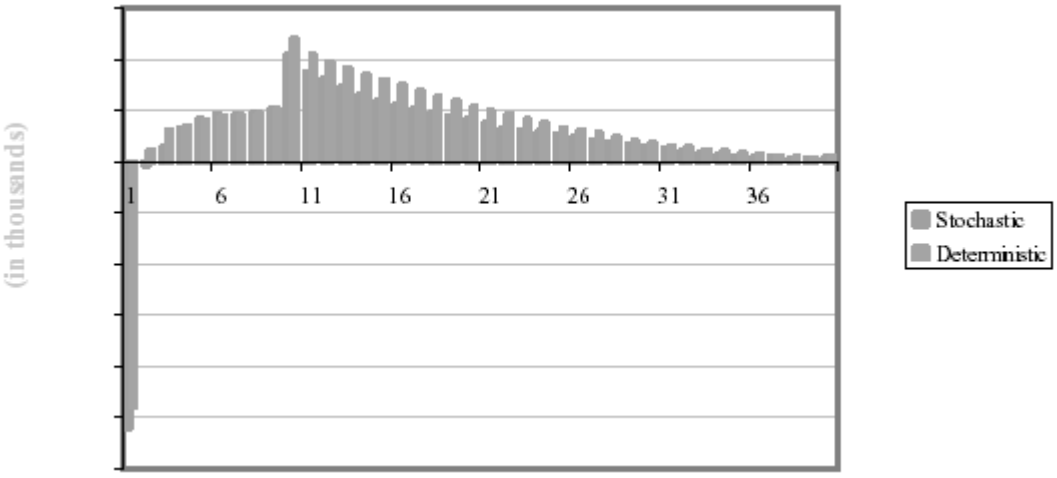

I would agree with what Hubert was saying that you do need to be careful with VAs because sometimes the equity is so low that you can get numbers that are pretty volatile. I would also say that you can—in response to your question—look at either the average ROI or average distributable earnings and translate that into an ROI measure. There are various ways, but you should definitely reflect risk in your thinking and not just look at the median.

Let's recap the results on a real-world basis. The stochastic valuation of the guarantee took the EV at the end of 2004 down from $14.2 million to $10.0 million, and it took the value of new business sold in 2005 down from $1.4 million to $900,000. It took the return on new business down from 19 percent to 14 percent.

With that, I will turn it back over to Maria, who will talk about the risk-neutral valuation.

MS. TORRES-JORDA: Now we're going to see what happens when we perform an MCEV valuation. To be able to do this, the first thing that we have to do is generate risk-neutral scenarios for equity returns. The volatility that we assumed for the scenarios was very similar to the volatility that was assumed for the realistic stochastic scenarios, i.e., around 16 percent. But for these scenarios, the expected return is the risk-free rate. If we compare Chart 7 to Chart 4 (that Mike discussed earlier), we can see that now the guarantee is going to be in the money in more than half of the scenarios. In Chart 4, we saw that it was going to be in the money in around 25 percent of the scenarios. Since the expected return under a risk-neutral world is the risk-free rate, the numbers of scenarios in which the guarantee is going to be in the money is higher than in a realistic stochastic valuation.

Chart 7

Risk-Neutral Scenarios
GMDB Benefits Capped at 200%

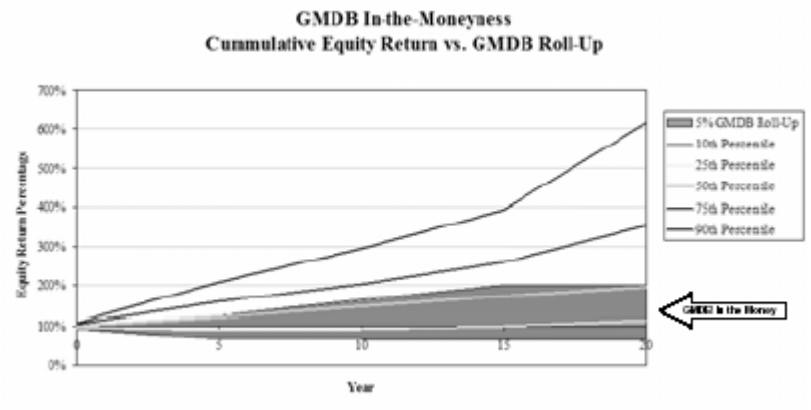

Let's go now to the results shown in Chart 8. The MCEV is going to be equal to the market value of the assets less the market value of the liabilities, which is basically the present value of the future liability cash flows discounted at the risk-free rate. When we projected future liability cash flows, since here we are doing a market-consistent valuation using the certainty-equivalent method, we assumed that we were earning the risk-neutral equity returns. We also have to take into consideration that there are frictional capital costs. In our example, we assumed that they were around 1 percent of the required capital. The EV as of December 31, 2004, for the in-force business is $2.5 million. For the value of new business, we can see that we have a slight loss.

Chart 8
Stochastic EV Results – Risk Neutral

Value of Existing Business at 12/31/04	
Market Value of Assets	$477,954,799
Market Value of Liability	(475,205,238)
Frictional Capital Costs	(240,764)
Embedded Value at 12/31/04	$ 2,508,797
Value of Business Sold in 2005	($ 574,048)

Chart 9 compares the results of the real-world stochastic method to the MCEV method. Starting from the first results that we showed, the deterministic EV was $14.2 million. Then we computed the real-world stochastic EV, in which we calculated the value of options and guarantees using realistic stochastic scenarios.

The EV of the business for the in-force block went down $4.2 million to $10.0 million. When we applied the MCEV methodology, the value of the in-force block went down to $2.5 million. The main reason for this is the difference between computing the value of options and guarantees using realistic scenarios versus using a risk-neutral valuation. The value of new business is also significantly lower under the risk-neutral framework.

Chart 9

Comparison of Results

	Embedded Value (Millions)	
	Existing Business At 12/31/04	2005 Sales
Deterministic Embedded Value	$14.2	$1.4
Real World Stochastic Embedded Value Real World Value of GMDB	$10.0	$0.9
Risk Neutral Stochastic Embedded Value Value of GMDB (Risk Neutral in Excess of Real World) Other (Discount rate, Frictional Capital, Cost of Capital)	$ 2.5	($0.6)

Theoretically, the top-down and the bottom-up approaches for calculating EV should produce similar results. In practice, this is not the case due to the issues around setting the appropriate risk discount rate. The question is: what top-down discount rate would produce equivalent results? This question is difficult to answer without directly computing the EV on a market-consistent basis.

The examples that we have shown don't reflect policyholder behavior. If we were reflecting dynamic behavior, such as reducing the lapse rate when the GMDB is significantly in the money, this would increase the cost of options and guarantees for both the realistic and the risk-neutral valuation. However, the increase will be much more significant for the risk-neutral valuation when compared to the realistic framework.

Let's see what the impact is going to be on other products. The impact depends on the valuation basis that we're using, i.e., whether we are doing a real-world stochastic EV or an MCEV, and it also depends on the product type. What you would see in general is that the value of options and guarantees computed using a risk-neutral valuation would tend to be higher than the value of options and guarantees calculated using realistic stochastic scenarios. If we leave the value of options and guarantees aside, one of the key differences is that products with a significant investment component, such as a single premium deferred annuity (SPDA) or a universal life product, would tend to show a much lower value under MCEV, because under this kind of framework you cannot capitalize future investment spreads.

MR. HUGHES: We did put together a practitioner's guide, if you will, on stochastic EV. If you were going to head down a stochastic EV path, my advice would be, don't

underestimate the effort. Many of us may have some comfort level with EV and some of us may have done some stochastic analysis, but I think stochastic EV does take things to the next level.

If I were to group recommendations, one bucket would be to lay the right foundation. By having the right foundation, I mean get the right team: people with EV expertise, people with the capital markets and the finance theory experience, people with the modeling expertise and people with the technology expertise. Make sure you have the right framework, because you don't want to take the old "ready, shoot, aim" approach. It's worthwhile to put some thought into the framework that you're going to be following. Get the right modeling tools. The traditional tools that we've been using might not be suitable for where you're going to want to be for this purpose. Typically, the scenario generators are a big issue. I would encourage you to make sure that you have the right scenario generators and understand what's embedded in them before you go too far. Get the right computing environment to get the runtimes that you need. Obviously, you're also going to need the stochastic models.

The other point I'd make is that it's very important to start simple and build. That's my other broad recommendation. If you try to throw everything in the pot at once and then get comfortable with the results, you're never going to get there. You're never going to know if the results are coming from a bug in the model that you haven't found yet or if it's something real with the business. I would say start simple, get comfortable with the models on a simple level and then, as you add complexity, get comfortable with each new level of complexity that you're adding. You'll be much more likely to end up with a good result and not end up going down a blind alley and then having to restart.

In terms of closing remarks, I think we are in a new world where stochastic valuation is upon us. We're seeing it with stochastic EV. We're seeing it with C-3 Phase II, economic capital, proposed statutory reserving guidelines, Standard of Practice (SOP) 03-1 GAAP reserving guidance, etc. We're entering a new world as actuaries. I have every confidence that we'll be able to get there, but we're going to need to elevate our game. We're going to have to brush off some of our finance theory textbooks, get the right tools and technology and put some emphasis into our quality assurance, because it's easy to run off the rails here.

MR. MUELLER: I'd like to make two comments on the assumptions before the next presentation. I think we said "separate account growth rate" assumption. We probably want to be specific and say that if your separate account includes a proportion of the assets in bonds, you probably want to take that into account. It shouldn't be just the equity return assumption. It should be a weighted return assumption if your separate account has bonds. In many cases, even if you don't have a fixed option, you're going to have some of the assets in bonds.

The other comment is that when we talk about spreads, I think we need to be

careful to distinguish between a risk-free rate and a Treasury bond rate. For example, it would not be necessary to take out all the spread from a corporate bond all the way down to a Treasury bond. You probably want to have the equivalent of a swap rate in your projection because that's a risk-free rate. You just have a liquidity premium that you pay for holding a Treasury bond, so in a sense the equivalent risk-free rate might be 30, 40 or 50 bps above the Treasury rate. You don't need to take your spread all the way down to the Treasury return, but you do want to take it down to a risk-free rate. There was a very good paper written by the Academy about two years ago on fair-value approaches, which did talk about how you determine the risk-free rate from Treasury rates.

Last, but not least, we have somebody from the other side of the pond providing a perspective on stochastic EV and MCEV. We're pleased to have Penny Coulthard with us today, who flew in just for this meeting to do this presentation. She's a senior consultant in Tillinghast's London office on the life insurance side. She has been with us in London since December 1998 and spends most of her time calculating market-consistent and traditional EV. She also works in mergers and acquisitions, restructuring, realistic balance sheets and individual capital assessments. She is a Fellow of the Institute of Actuaries in the United Kingdom.

MS. PENNY COULTHARD: I've been at some good sessions today, and I've been struck by how similar the issues that you're facing at the moment are to the issues that we're discussing in the U.K. actuarial profession at the moment, whether it be risk management, stochastic modeling and how to make sense of these different models and different results, or even how different people are viewing the actuarial profession.

Most of the work that I've done in the United Kingdom over the last three to four years has used stochastic techniques to come up with solutions. Gone are the days when we were able to do deterministic solutions. I'd love to say that that's as a result of the U.K. actuarial profession pushing us forward and making sure that we strive for best practice, but I'm afraid it's actually being pushed by the regulator in the United Kingdom.

Following the demise of Equitable Life, which led to the Morris review, the regulator transformed the way that we actually prove that we're solvent in the United Kingdom. For starters, for certain participating business, we now have to do a market-consistent balance sheet to prove that we're solvent, using risk-neutral methods to come up with a market-consistent value of our liabilities. That's now being published by companies, as of the end of 2004.

Further, we also need to calculate something called an individual capital assessment. This is a company specific economic capital calculation—we have to prove that there is only one event in 200 where we go insolvent over a year. This involves doing do real-world stochastic projections for a year with risk-neutral assessments of capital requirements thereafter.

This has been a lot of work—but it is now becoming business as usual—and as such management are looking to the models and tools that they have invested so much time and effort in building to gain a better insight into the risks and values within the complex businesses that they are running. MCEV is one such application of how firms are moving from a period of regulatory compliance to one of business insight.

I am going to spend the next 30-40 minutes talking about three case studies involving MCEVs, drawing out some of the issues that we come across when we do these calculations.

As background, EV has been widely used for a number of years as supplementary information by companies in their Report and Accounts. However, these EVs generated much criticism from the analyst community, primarily for the lack of consistency and transparency between companies and also for the lack of allowance for options and guarantees. To overcome this, the CFOs of the biggest European companies got together and created the “EEV Principles” which aim to create the consistent of the methods applied to calculate the EV and to increase transparency through additional disclosure. EEV results are now being published thick and fast. Since I submitted the slides, AEGON has also published EEV and a number more are due shortly.

Now the primary reason for these companies publishing on EEV is the adoption of new external reporting standards. If we look now to MCEV, a number of companies have published these over the last two years.

AMP published MCEV in its explanatory memorandum when it demerged its Australian and UK companies. R&SA then published MCEV to support capital raising. There have been a large number of additional cases when MCEV has been used in M&A—and this has become a must-have in these transactions. HHG (formerly AMP UK) then published their second EV during 2004 as end-year supplementary reporting. RAS has used MCEV techniques in their EEV calculations and we are aware of a number of other companies who have not published yet, but who are looking to adopt MCEV. The majority view is that most companies will use MCEV in EEV in the future.

MCEV is, therefore, providing a framework for pulling together all these different things about which managements are concerned: a mark-to-market basis for managing the business; communicating with shareholders on the value and the performance of the company; pricing mismatching risks, guarantees and options; balancing risk and reward, evaluating de-risking strategies (that's particularly pertinent in Europe, given falling interest rates in equity markets); assessing value in a restructuring or a merger-and-acquisition (M&A) situation and also, as I said, regulatory reporting is now on an MCEV basis.

My first case study thinks about movement analysis in an MCEV. Now, as Hubert said, whilst an EV can be quite an interesting number, for example, in an M&A

situation when you're trying to assess the value that you should place on a company, it's actually the movement in the EV from one period to the next that allows the management and the shareholders of the company to see how well that company has performed. The movement analysis then takes that further and splits the movement down into different chunks to show you from where the profits have been coming. This allows the management and shareholders to see which bits of the company are doing well and which bits are doing less well.

The traditional EV techniques would smooth the profits somewhat through the margins that were in most of the assumptions. We don't get that in an MCEV. It gives a clearer picture as to what has really happened over the period.

HHG were publishing their second MCEV, and the company needed an approach to communicate the movement in a meaningful way. They were the first company to do this type of movement analysis, and so they were coming at this from a blank sheet of paper.

Under an MCEV framework, we can think of a life insurance company as comprising two distinct businesses: an insurance business and an investment business. The insurance business takes on all of the insurance risks in the company: persistency, mortality, expenses and operational risks. They all sit within that insurance business. It does not contain any investment risks. It's assessed on the basis that that company is invested in a replicating portfolio of assets, so there are not going to be any investment profits or losses emerging within that business. Those investment profits or losses are going to fall within the investment business' balance sheet.

It's very unlikely that a life company will be invested in a replicating portfolio of assets. For example, it may not be possible to get put options on real estate or you might not be able to get long-enough dated bonds. Even if you're trying to get into a hedge position, you probably are not going to manage it perfectly. Some companies may not be trying at all and may be taking a mismatched position with a view that they're going to be generating profits through taking on a slightly riskier investments stance. That mismatching risk will be coming through in your investment business.

If we think of insurance companies as comprising these two distinct businesses, an insurance business and an investment business, then when we think about the profits that emerge, we can think about the profits that emerge in each of these two distinct businesses. Within the insurance business, the in-force business will generate profits if there are variances against assumptions. This might happen because you better your best-estimate assumptions so you have an experience variance, or you might actually change your best-estimate assumptions. For example, you might get some better mortality improvement factors that make you reassess your basis as to what you think is going to happen in the future.

Within MCEVs we assume that insurance risks are diversifiable and as such we include insurance assumptions at best-estimate. Therefore we are not expecting to see any prudence margins being released from one period to the next, but will only see insurance variances arising from experience variances or assumption changes.

The new business will also come in here. At the point of sale, new business will generate value if the market-consistent value of the future income exceeds the market-consistent value of future outgo, including the cost of capital. Between the point of sale and the valuation dates, that new business will then generate investments and insurance profits as for the in-force block.

Let's look now to the investment business. As I said, it's very unlikely that a company will genuinely be invested in a replicating portfolio of assets. As such, as economic conditions change over time, we'll see that the assets move differently to the liabilities following a change in economic conditions. This will generate a mismatching profit. In addition, the free assets, or the assets not backing the market-consistent liabilities, will also grow with investment returns.

HHG decided to take this view of their business. Chart 10 is the movement analysis that they published in their December 2003 Report and Accounts. The MCEV had gone up £371m over a six-month period. They had 21 million pounds worth of opening adjustments. They then expected a return of 34 million. This expected return will have come purely within their investment business because, as I said, in an insurance business they weren't expecting any profits because everything had been put in at best estimates. They then created some value through their investment management and insurance management. They also had some capital movements (this was just after the demerger; things were still shifting around quite a bit). And finally there was an "Other," or unexplained movement, of 14.

Chart 10

Case Study 1
HHG's MCEV Analysis of Movement

	£m
Consolidated MCEV (excluding explicit agency costs) at 30 June 2003	900
Opening Adjustments	21
Expected Return	34
Value created through Investment Management	143
Value created through Insurance Management	92
Capital movements	67
Other	14
Consolidated MCEV (excluding explicit agency costs) at 31 December 2003	1,271

This gave analysts the reason why the EV had moved so much over the period. Also, I think that unexplained movement is quite interesting. An analysis of movement is a requirement under the Embedded Values Principles, but it is also a good check on the numbers. If you have done a detailed and rigorous movement analysis, you should gain quite a lot of comfort that you've not introduced errors into your models or processes between the two valuation dates.

My second case study is around setting the risk discount rates. We've already had a talk on this, but I think that it's so fundamental that I'm going to briefly go over it again. The question as to how you set a risk discount rate in an EV has been one facing the actuarial profession for a long time now, and we've finally come up with an answer with which we're happy. Unfortunately, we've actually come up with two answers. So RAS wanted to decide the best approach to setting its risk discount rates in its EV calculations. It had two possibilities: the top-down approach and the bottom-up approach. It wants to decide which one to use, so it needs to think about the pros and cons.

The first option was this top-down approach, where the risk discount rate is based on WACC or CAPM methodology with one single risk premium and the cost of options deducted from the total in-force values. There are a number of benefits to this approach. It's widely used by their competitors. At a time when RAS was putting together these numbers, all of its competitors had used this approach. For RAS to do something different would be stepping away from the crowd. However, we are aware of a number of companies that are going to be using a bottom-up approach and just haven't published yet. The risks are partially captured in the discount rate and partially in the projected inflows, so that's more akin to what they were used to under their old EV reporting. There's an overlapping of stochastic and deterministic frameworks, which, again, is more familiar.

However, there were some serious downsides to this approach. As I think everybody has mentioned now, the discount rate is not differentiated by product. You have one discount rate for all of your products, irrespective of the riskiness in those products. Now think about your value of new business. Suppose you've sold lots of risky business in the past with lots of guarantees in it. That means you're going to have quite a high discount rate coming through in your WACC calculation. However, suppose now you don't want to be selling that business anymore and you're selling something without those guarantees in it. You'll be using too high a risk discount rate to value your new business and, as such, you're going to be pushing down your value of new business, which may not be something that you want to be doing.

There's a lot of subjectivity in setting top-down risk discount rates. If you actually go through one of these exercises, you start to realize how many judgment calls you're making. From where should I get my beta? How many years should I average over? What equity risk premium should I assume? How do I adjust the beta

to remove the general insurance and my banking parts of the business? There's a lot of subjectivity that goes into those decisions. When you start to go through it, you realize that there's a large range of possible outcomes. In addition, you are basing the risk discount rate on market informational betas, for which there are a number of well-known problems. Finally, the stand-alone cost of options may not be meaningful if it's not consistent with markets.

Chart 11 shows betas for 30 sample insurance companies at the end of February 2005. You can see that the betas do vary quite significantly, from about 0.5 to 1.7. In some cases, it's quite clear why those betas are varying so much. For example, some of the U.K. with-profits companies have high betas to reflect the fact that the markets are quite uncertain about the risks involved in with-profits at the time. However, in other cases, it's less clear why the betas are as they are. There seems to be quite a lot of noise within that sample.

Chart 11

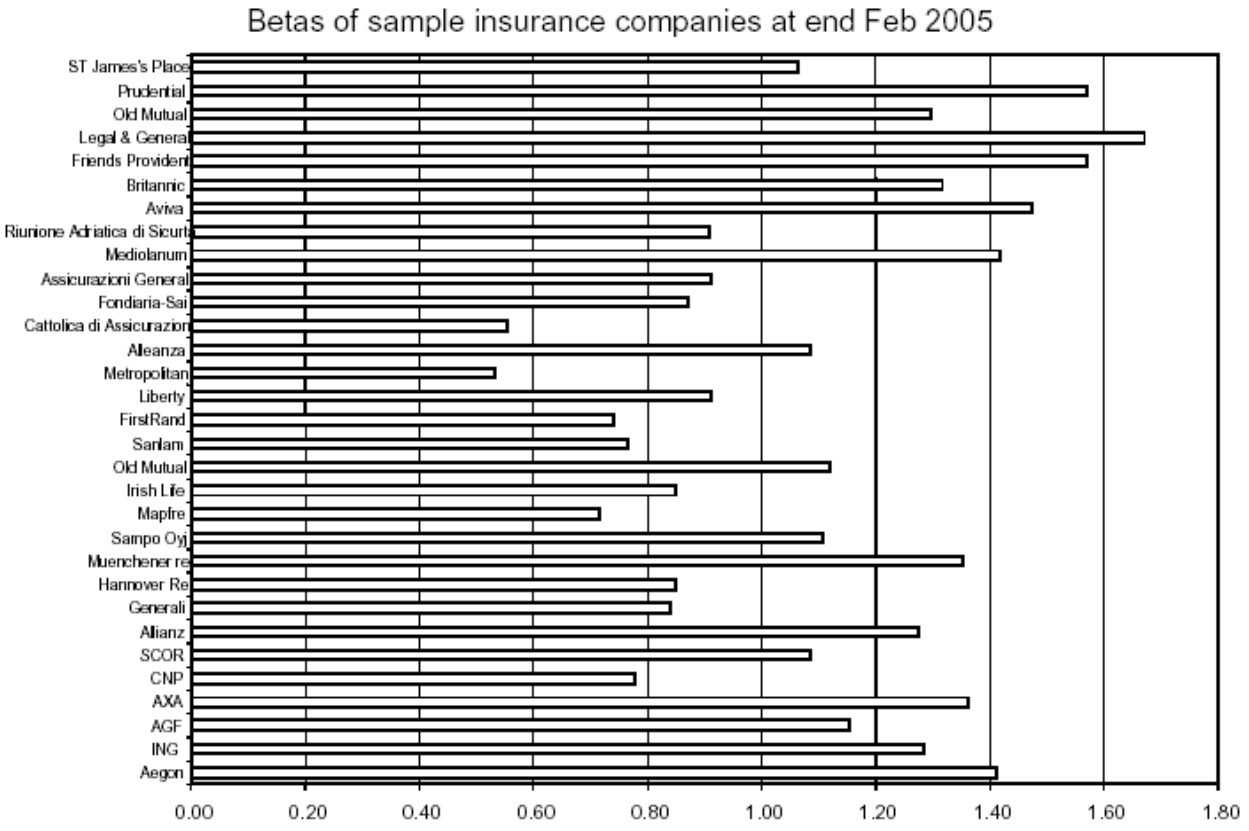

The second option RAS considered was a bottom-up MCEV approach. In this approach you allow for the cost of market risks in the valuation of the cash flows.

I'm going to go over the example in Chart 12 because it's so fundamental to what we're trying to do in this type of valuation.

Chart 12

Case Study 2
Option 2: Bottom-up MCEV Approach

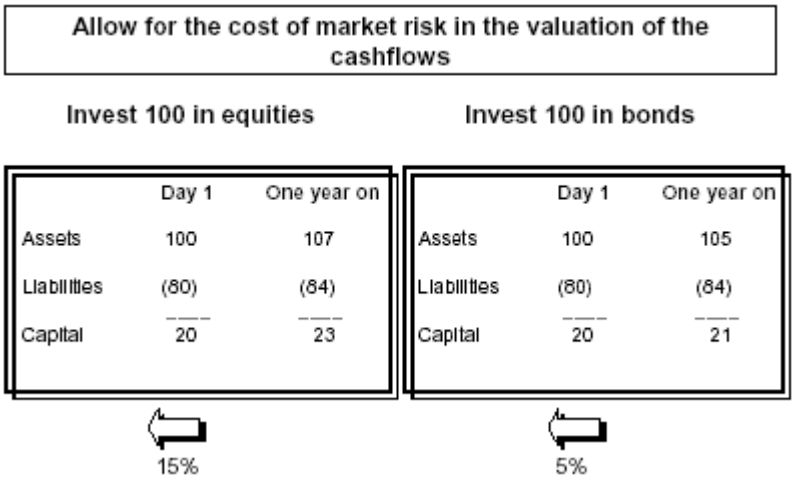

Say we invest 100 in equities. On day 1, we have 100 of equities and fixed liabilities of 80, so we have capital of 20. One year later, we expect our equities to grow at 7 percent, 2 percent above the risk-free rate, to 107. We expect our liabilities to grow at the risk-free rate, because they're guaranteed, we know with certainty, to 84, giving us capital of 23. We know that the market value of that capital today is 20. If we solve for the risk discount rate, we get a risk discount rate of 15 percent. This is quite a high risk discount rate, reflecting the fact that this is really quite a leveraged transaction. You're backing a fixed liability with equities and the risk discount rate is picking up that mismatching risk.

If, on the other hand, we invest 100 in bonds, then our assets grow at the risk-free rate. We expect them to grow at the risk-free rate to 105. Our liabilities also grow at the risk-free rate, so our capital grows from 20 to 21 and we can solve for a risk discount rate of 5 percent. The lower discount rate is reflecting the fact that we're not taking any mismatching risk within this transaction. So within a market-consistent valuation framework, we'd be valuing the equities example at 15 percent and the bonds example at 5 percent. In a top-down type of case, we'd be valuing both of them around 10 percent; it would be overvaluing the investment in equities and undervaluing the investments in bonds.

RAS' second option was to calculate bottom-up risk discount rates based on the valuation of the risk factors with differentiated risk premia by line of business and country. A benefit to this approach, as we said, is that the discount rate is tailored

to the specific risk factors of the different products. This has an important impact on the risk management because you can see that the products with a high risk discount rate are riskier. There's also a consistent link between the stochastic and deterministic framework.

I'm not going to pretend that there aren't some downsides to this. There are more calculations. It's quite difficult. We're quite lucky in the United Kingdom with having this stochastic framework forced on us by the regulator. We have models in place to do all this stuff already. But there are quite complex calculations, as we were saying earlier. There are new explanations required in the disclosures, particularly for RAS, when they were the first people to go out to market with this approach. However, RAS concluded that this was the best approach, and this is what they've used to set their risk discount rates in their EEV calculations.

Let's look at the third and final case study. Moving to an MCEV framework will make some of your products look more profitable. I think that's important to say. It won't decrease the profitability of every product line. Some will look better, but some other ones will look worse. We're going to have risk discount rates that value by product and reflect the riskiness of the underlying products.

Our third case study considers a company for whom one of their territories sold credit spread products. They were concerned that moving from their traditional EV approach to an MCEV would wipe out the value of the subsidiary. Why would that happen? Under their traditional EV approach, they used a WACC approach to set their risk discount rate, so there was one risk discount rate across the entire group. They then calculated the EV by projecting forward the cash flows under that product, including a risk premia that they expected to get on that credit business, and then discounting it back. They were effectively capitalizing that risk premia within the value.

Let's go on to the MCEV. Although the EV of the total group wouldn't actually change (as I said, top down equals bottom up, in theory), the product would be discounted at a higher rate to reflect the riskiness of the corporate bond cash flows. Under an MCEV, the MCEV does not bank these credit risk premia, but rather the credit risk premia come through in the EV profits only once the risks have been borne.

The company saw a number of problems in moving to an MCEV basis. In particular, if they really believed the MCEV was more appropriate, then there were implications for their new business and pricing strategies. It was quite a big step for this company to go to a market-consistent basis for pricing new business. Why? Well, their current pricing basis capitalized these credit risk premia upfront. Their competitors were also passing on risk premia before they had been released from the risk, so it was difficult to see how they could compete if everybody else was taking this approach.

On the other hand, if you think that your profits are coming from credit risk gains, why not simply borrow the money and invest in corporate bonds rather than take on all the extra risks of investing via an insurance company? This company did eventually decide that it was going to move to an MCEV framework, and this did have some implications for the way that it was managing this block of business. They're planning to tighten their pricing basis. They've accepted the fact that they'll only see those investment profits from credit risk premium coming through when that risk is being taken. They've tried to find ways to transfer some of that credit risk to policyholders. They've also tightened the risk management on their credit book through tighter asset-liability matching.

In conclusion, just like the radical long hair of the Beatles stirred up fear among traditional audiences, so does the market-consistent message strike fear among U.S. actuaries and GAAP fans (it's fair to say the United Kingdom as well, to be honest). Why is it frightening? There is some complicated modeling to do, but then, again, we're actuaries and we love really complicated modeling. We're scared that the value of the business is going to plummet. Again, I would say that some products come out worse, but some products do come out better. We're concerned about the volatility of the MCEV. Again, I think that's fair. We've taken out the smoothing that was being covered in the traditional framework. However, that volatility is real. It's in the business and so, perhaps, we should be recognizing that. There have been some benefits from implementing MCEVs in Europe. We now have a more rigorous framework for calculating EVs. We have numbers that the capital markets understand, and there is an objective recognition and pricing of risk.

MR. MUELLER: We will now take questions from the floor.

MR. JAMES SCOTT RUSSELL: When you're capturing the price of the options, you're running the in-force model deterministically and then subtracting the stochastic result. Does the in-force, though, have GMDBs that are in the money at the model start date, or are those wiped out completely? There are costs of options from the past and then there are costs of options in the future. What are we trying to measure there—all of the past and future options, or just the future?

MR. HUGHES: The option value has two components. It has the intrinsic value and the time value. The traditional EV captures the intrinsic value, the value based on a projection starting from where you're at today. The stochastic methods capture the time value of that guarantee. There's a partial recognition of the guarantee in the traditional deterministic to the extent that it's in the money. The stochastic captures the rest of it.

MR. RUSSELL: So the deterministic implicitly has the in-the-moneyness.

MR. MUELLER: Yes. Even a traditional EV would include a costing of the options and guarantees on a stochastic basis, but plugged into a single scenario. In other words, you would do a side calculation of your GMDB cost using a stochastic model,

maybe come up with 15 bps and then plug that cost back into your traditional EV model. Otherwise, in most circumstances, you would get no cost at all for options and guarantees in a traditional EV model. So you do a side calculation where you calculate the cost stochastically, and then you plug that into your one single scenario as an annual cost in bps, but the bottom line is that it does not capture the full cost of the options.

MR. RUSSELL: But the cost that you're capturing is only the future.

MR. MUELLER: That's correct. It's going forward from the valuation date. But essentially the cost is determined by where you're starting from. To the extent that you're in the money, you're going to have a higher cost, because if you did a separate stochastic model where you're 20 percent under water, and the stochastic model comes out with the cost, it's going to be much higher than if you're 20 percent out of the money.

MR. RUSSELL: Can you speak to how risk-neutral scenarios are generated or calculated? In particular, the value of new business sold in 2005 under the risk-neutral EV is now negative. What happened there?

MS. TORRES-JORDA: You have to be very careful when you generate the risk-neutral scenarios that they reproduce market prices. Once you generate the scenarios, you need to check that if you want to price an instrument that is being traded in the market, you will reproduce the market price of that instrument as of the valuation date. If that's the case, then your risk scenarios have been generated correctly. That's one of the things that you need to do to make sure that you are doing the right thing when you're implementing a stochastic framework. You definitely have to do that whenever you're generating risk-neutral scenarios. Make sure that you reproduce market prices.

Regarding what you were saying about the value of new business, what is happening there is that under an MCEV framework, the value of the options and guarantees are significantly higher. They're the true value of the options and guarantees. Doing a realistic stochastic valuation for options and guarantees will never give you the right values. It's not consistent with how markets price those kinds of options. In this particular example that we were showing for the value of new business, the value of the options and guarantees under the risk-neutral framework were significantly higher than in the realistic framework, giving you a negative value of new business.

MS. COULTHARD: Yes, that's absolutely right. Basically, what we're doing is looking at our life insurance liabilities and saying what sort of asset it is like. A guarantee might be like, say, an equity put option. We're trying to find a way to value that equity put option in line with how the market would have valued them. We're using stochastic techniques to value those equity put options, so we check that our stochastic scenarios, if we use them to value equity put options give a

result close to the observed market value.

In terms of what they look like, I think all of the asset classes on average will be return the risk-free return, but, as Maria said, it's just a tool for coming up with the market-consistent valuation. The volatility within the scenarios, say the equities, will be higher than the volatility of, say, the bond returns, to reflect the riskiness in those equity products and the fact that your guarantees are more likely to be uncovered if you're invested in equities rather than in bonds.

MR. MUELLER: The whole issue of validating your risk-neutral scenarios to current market prices is something that investment banks do on a daily basis. This is not in any way a new science. It's just that we, as actuaries, are coming onto that a little later than the investment banks. That's why I was talking earlier about us needing to be careful about market arbitrage. If you're looking at a block of business, which, let's assume, has certain options and guarantees (whether it's credit spread or guarantees on investment products), and you're valuing that block using traditional actuarial pricing techniques, if you looked at the examples that we were just doing, you're overpaying. The seller will be very happy and will not tell you that you're wrong. But, as a buyer, you're overpaying.

I have seen transactions where the buyers were willing to pay traditional actuarial prices until we told them that the options were not valued on a market-consistent basis. In at least one case of a transaction in Switzerland, it was a large portfolio of unit-linked-type products that had investment guarantees, and the seller wanted something like 2 billion Swiss francs. On a market-consistent basis, the value became minus 100 million Swiss francs, so it was a huge difference. The buyer would have been willing to pay the price if they hadn't had an external advisor, because they thought it was just the way to do it. You need to be careful in determining the right market price.

MS. JULIA LYNN WIRCH: Was the roll-up guarantee that you had only electable after 20 years?

MR. HUGHES: No, it was applicable from policy issue.

MS. WIRCH: From the way I looked at it, it looked like after 20 years it covered the 25th percentile barrier for the real world, as shown in Chart 4, or the 50th percentile barrier for the risk neutral, as shown in Chart 7.

MR. HUGHES: Yes. It's actually a little bit more likely. It's more like 25 or 30 percent.

MS. WIRCH: Yes. One of the things that's evident with this type of analysis is that when you're comparing an equity return, which is quite volatile, and a guarantee, which is 3.75 percent or something like that here, even though it's a fairly small guarantee, within the first five or 10 years, it's very risky. What you see in Chart 4

is that in year three (or four, maybe), it almost approaches the 50th percentile.

MR. MUELLER: Right.

MS. WIRCH: If you don't put a block in saying that if you take your money out before five years (or whatever) then you don't get that guarantee, then you're giving yourself a lot more risk than if you do put that block.

MR. HUGHES: I think that's a very good point.

MS. WIRCH: Chart 7 shows it a lot better. It's a little more evident, but it's not until year 15 that it's higher than the 50th percentile. But usually we think about 20 years as the shortest duration.

MR. MUELLER: In the second economic capital session this morning, the example put up on the guaranteed minimum withdrawal benefit (GMWB) case study assumed a 1995 issue, single issue year. The equity market in the second half of the 1990s was very bullish, so the market went up, the GMWB had a reset basically on the market peak and then people started taking withdrawals, which is almost our worst-case scenario. That's why the results looked as bad as they did. It was a 1995 issue reset at 2000 just at the peak of the market, and then people started taking withdrawals.

MS. BEVERLY ELIZABETH STEINHOFF: Penny, you were talking about how MCEV is a volatile measurement tool. If you're using it to manage your business and to make strategic decisions, what are some of the ways companies are starting to control the volatility of MCEV?

MS. COULTHARD: One of the main areas where the volatility arises is from the asset-liability mismatches in the business. If you are invested in a portfolio of replicating assets and you invested the rest of your assets in bonds, you wouldn't see this volatility coming through. Companies that can't afford to have this volatility and decide they don't want that volatility will be looking to de-risk their asset-liability mismatches through hedging strategies. That's happening an awful lot in the United Kingdom at the moment.

In terms of the other risks, like the expenses risk, many companies are outsourcing their administration at the moment so they're not seeing expense variances coming through. They've locked into a fixed expense amount per annum in which it goes up at a fixed inflation amount, so that risk is going off balance sheet now. Persistency is a bit more difficult to manage. Another thing that companies are occasionally doing is securitizing bits of their business that they feel are going to be particularly volatile, where they don't have the experience or capital to manage it. There's an awful lot of de-risking going on in the United Kingdom.

MR. MUELLER: You could add mortality risks through reinsurance, right?

MS. COULTHARD: Yes, absolutely.

MR. MUELLER: The focus is on taking the volatility out of the business as much as possible. In the United States, we see companies increasing their retention limits: buying less reinsurance, buying more stop-loss, buying more coverage at higher levels. In Europe, we see companies actually buying more reinsurance of the traditional kind because they want to take out some of the volatility of the business when they look at it on a market-consistent basis. That is an interesting dichotomy between the two markets.

MR. GUY HORTON: Risk-free rates: Europe, not so hard; United States at the moment, not so hard but in the future, maybe not so easy. But Brazil and Argentina, am I a foreign investor or am I a local investor, because the risk-free rate varies greatly if you're assuming some sovereign risk or not. In Brazil, for example, two years ago, we had maybe a 12 percent credit spread. If I'm a local investor, I can use the put option, because if the government defaults, then probably as a company I can sort of ride on the back of that. If I'm a foreign investor, I cannot necessarily do that. How do you do your risk-free rates in those volatile situations?

MS. TORRES-JORDA: In markets like Brazil and Argentina, coming up with a risk-free rate is a difficult task, because you cannot really look at the government bonds to come up with a risk-free rate; it is not risk-free. The issue there is much more difficult than in countries like the United States or the United Kingdom, where you could look at it from the point of view of the local investor. When companies consolidate all their different MCEV values for different countries, they would just convert different values from different countries based on these bonds' exchange rate as of the valuation date.

MR. MUELLER: I have seen valuations being done both ways. If you're representing a local company, they tend to look at it with more of a long-term view toward the market, tend to be more willing to pay for some of that spread if the spread is there and maybe assume it stays there, or they have their own history of the market. If you're talking to foreign investors looking at a certain market, whether they go in for the first time or they expand where they already are, they tend to look at it the same way they would look in every other market, and they would tend to discount those spreads as less believable.

When I worked in one of our European offices, we worked in the Polish market a lot when it was clear that Poland was going to be becoming part of the European Union. Even then, in the late 1990s, it was clear that eventually, in the next five years, Poland would become part of the European Union. They had interest rate levels greater than 20 percent and huge corporate bond spreads, but nobody was willing to project those spreads forward for 10 or 15 or 20 years. You basically had sort of a timeline, where you knew the rates were coming down and the spreads were coming down, and all the valuations that I saw in that market were done on

that basis. It's a bit of a judgment call, but, certainly, I find that the local investors tend to be more subjective than the foreign investors looking at a particular market.