2003 VALUATION ACTUARY SYMPOSIUM
September 11–12, 2003
San Diego, California

Session 35TS
Hedging and Other Mitigation Techniques

Instructors: GEORGE CHRISTOPHER
MICHAEL J. O'CONNOR

Summary: This session describes current risk mitigation techniques and their impact upon valuation systems. Topics include: hedging mechanisms, policy updates, regulatory relief, anticipating evolution of GAAP measurement, product "quarantine," realistic assumption changes and maximizing return and minimizing litigation.

MR. MICHAEL J. O'CONNOR: There are two panelists today, George Christopher, from Milliman USA, and myself, Mike O'Connor, from Tillinghast. George has a Ph.D. in algorithms and a lot of other things. He's in the financial risk management practice at Milliman USA in the Chicago office, and he recently got his FSA. I'm an FSA, and I'm in the Minneapolis office of Tillinghast. I'm also involved in risk management issues, asset/liability modeling (ALM), hedging and a lot of other things relating to fixed and variable annuities (VAs), for example. The focus of the presentations today will be on the variable annuity side.

George is going to give a case study of a hedging technique, and I'm going to cover some other questions and issues around some possible potential mitigation techniques for VAs.

There are several topics that I'm going to discuss, especially for the inforce business. If you try to manage some of your risks beyond using a hedging strategy, you have to determine where your risks are, what product features you need, potentially fund allocations and that type of thing. I'm going to go through some case studies using the C-3 Phase II capital framework as a way to determine the relative risks from different products, different benefits and different dynamic lapse assumptions, for example. Then I'll talk about some potential product design

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

changes, assuming these changes would be more applicable to new business, but there would be some potential benefits that you could add to your inforce business to come up with a different risk profile.

I'll talk about product balancing and then give a reinsurance example, not to necessarily affect the underlying cash flows, but as a way to perhaps manage your reserves and capital position. I'll make a brief mention of some of the recent GAAP developments. There are other sessions that are going to cover some of these issues in greater detail.

This is a quote from a high-level individual at one of the reinsurance companies: "The best risk management tools will not solve the problems created by poor product design." But we are where we are with respect to designs in our inforce, and so I'm going to try to address some of the issues about what to do with that.

Again, the question of determining and prioritizing risks, like in this C-3 Phase II framework, would be effective most likely at the end of next year. It's a true risk-based capital (RBC) framework. I'm hoping that most of you are familiar with it because I won't go into it in much detail. I'll show examples of some numbers, but I'm not going to go into the framework itself that much. It does allow for hedging offsets in determining capital. It's very similar to the Canadian framework in a lot of ways. I'll go through a case study to analyze some different risks for a typical guaranteed minimum death benefit (GMDB).

I'll give you the product specs. It's a plain vanilla product: newly issued VA to a 65-year-old, seven-year surrender charge, grading linearly from 7 percent down to zero maturity age of 90, and mortality and expense (M&E) of 165 basis points (bps). In the modeling we did assume that there was a guaranteed minimum income benefit (GMIB) on the policy as well; there's an investment fee of 75 bps, but 25 of that is a reimbursement back to the insurance company. So, the GMDB is the greater of paid premium, annual ratchet up to attained age 80, and the premium accumulated at 5 percent up to attained age 80. We mentioned the GMIB, although the modeling results we're going to be talking about are limited to the GMDB.

What was tested? This is new business only. The numbers I'll show you are new business only. Inforce would be different, if not dramatically different, depending on when the policies were issued and the specific history of those policies, especially if they have a ratchet design. We'll talk about static and dynamic lapses and how that can influence the capital requirements. I'll talk about partial withdrawals, including dollar-for-dollar. I'll talk about transfers to the general account and transfers to a less risky fund like a generic bond fund (people transfer out of an equity fund).

The base case result for this 65-year-old is that the C-3 capital charge would be 88 basis points. I'm about 99 percent sure that the way it will flow through your RBC

work papers is that that's the 100 percent factor. If your target were 225 percent, you'd have to multiply all these numbers by 225 percent, or 2.25.

Again, the base case is 88 basis points. We need to keep that in mind as we go through here. That assumed just some base lapses. Now we start including different lapse patterns. The first one is that the lapses increase in a down market, and obviously there's a slight reduction in the capital requirement. Then we did a couple of different dynamic lapse assumptions, where the lapses were reduced the more and more that the GMDB became in the money.

Before I proceed further, I should mention where this case study came from. Two individuals in our Atlanta office, Doug Robbins and Marc Altschull, recently presented a paper on stochastic modeling of policyholder behavior, and this is drawn from it. It's kind of a summary of some of the higher-level results of that paper that they presented.

Again, if the base is 88, if lapses increase in a down market, you get a slight reduction. But if you have a severe lapse reduction, you'll have to have an extra 25 bps, and, again, you need to multiply these numbers by whatever RBC ratio you target. If there's a milder lapse reduction, if the GMDB is in the money (ITM), then it's an additional 20 basis points.

We also did some static lapse adjustments. What's interesting is that, if you double the base lapses after the surrender charge period is up, it doesn't affect the base numbers that much. One reason is that this calculation is done at issue. It's a seven-year surrender charge pattern, but it does not show a whole lot of sensitivity to what happens after the surrender charge. I believe that the shock lapse that is built into this is a 35 percent shock lapse after the surrender charge is done, and it's around 15 percent surrenders after that.

This one, and also the one where we tested transfers to a less risky fund like a generic bond fund, is trying to get at the question of whether that behavior increases risk to the insurance company or not. If it reduces risk, potentially through some bond fund opportunities, you might be able to encourage people to switch to a fund. Looking at just going to the general account, the capital increased by 33 basis points. It's a pretty simple dynamic transfer formula, so as the "in-the-moneyness" got to be above 10 percent, there's a transfer of one half of that. (What we mean by "in-the-moneyness" is the ratio of the account value over the GMDB. Let's say the account value is down to 90 percent of the GMDB value. That's 10 percent "in-the-moneyness," and that's the way we're using that phrase in this paper.) If the ratio of the fund value is at 80 percent, the difference above 10 percent is 10. So, there'd be a 5 percent transfer per year to the fixed account. But most of this increase in capital is due to the 3 percent minimum guarantee, which I know is a big issue for a lot of companies these days. If you took out that 3 percent guarantee and took it down to zero, then the capital increase would be 15 basis points, not 33. The general conclusion from this is that the transfers help to

insulate the policy from future losses. Maybe to some extent the policyholder would be locking in a loss for the insurance company, but also it helps to reduce, or cap, the future losses for the insurance company.

It's a similar thing when people transfer to a less risky fund. Say the generic fund was a large cap equity fund, and we said that people can transfer to a generic bond fund. The Lehman Aggregate Bond Fund would be an example. If there were no dynamic lapses, capital would be reduced by four to six basis points. That's almost in half because, again, the base case is 88. However, if you introduce dynamic lapses, the reduction in capital is much more modest. It's only 17 basis points. But, again, this does suggest that that type of behavior will reduce the risk to the insurance company and helps reduce the future losses.

For some companies, partial withdrawals are a much bigger issue. If you have a proportional reduction in GMDB, that's just like increasing lapses. It's a dollar-for-dollar reduction. That's potentially onerous. All these sensitivities are definitely driven by the specific dynamic lapse patterns we used. Perhaps you would come up with different results, but this gives you a good idea of the magnitude and definitely the direction of some of these behaviors. The increase in capital is 217 basis points, and, again, you need to multiply that by whatever your RBC ratio is. This is at point of sale, not for an inforce block that's already ITM 30 or 40 or 50 percent. It will definitely be driven by the policyholder behavior function, but this is an area where there's clearly a tremendous amount of exposure for companies.

How about higher issue age? For an 80-year-old, the death benefit is only return of premium by definition, but the base capital requirement is 195 instead of 88. It's a little bit more than doubling. There's clearly exposure by issue age, if you go from a 65-year-old to an age 80. The dollar-for-dollar provision is an extra 189. So, again, the total for an 80-year-old with a dollar-for-dollar provision is almost 4 percent, and, again, you have to multiply that by whatever RBC ratio you're targeting. Again, this is at issue.

Transfers to less risky funds do reduce the capital by a large amount. The reason we bring up these things is that this would be a tool to help determine your risks, prioritize where you might need to focus and then see if there are some potential solutions to help you deal with where you are. The approach was looking at the capital implications because the capital implications are looking at the tail events, and the C-3 Phase II will probably be effective end of next year. If anyone in the audience has a better handle on the timing, if it's a different timeframe than that, I'd appreciate your comments.

Once you've looked at risk in this framework, then you have to prioritize the riskiest benefits. Obviously the dollar-for-dollar is probably the riskiest. These are just GMDBs. The paper that Doug and Mark wrote also has a section looking at GMIBs and guaranteed minimum withdrawal benefits (GMWBs), with parallel sensitivities and transfers. If you have gone through this process, and you have figured out

where your risks are, then you need to explore some ways to modify existing policies. Frankly, the opportunity here is probably somewhat limited. But you can offer riders with offsetting risks, such as an enhanced earnings benefit on policies with a GMDB. I think you need to do the modeling to see how much that would help. There has been a fair amount of discussions at companies to do conversion programs. But from what I can see, there has been little activity in doing that.

If you wanted to do something like that, there would definitely be a whole range of issues that you would need to explore, such as making sure that you had an equitable program to offer people. Arguably, the worst thing that could happen is that you convert a financial risk into a litigation risk. Other issues would be your SEC filings and some possibility of causing incremental anti-selection. You can have potential changes to the funds available, like limiting the high-risk funds. In the inforce business you may not have much opportunity to do this, but definitely for new business, and potentially for some of your inforce policies, you might require asset allocation for some of the guarantees.

What's been happening in the market over the past year? I know this is talking more about the new business side. A lot of these benefits have been re-priced, modified or eliminated over the past year. The dollar-for-dollar benefit is probably one of the main ones where there's been a fair amount of modification or elimination. As most of you are aware, if not all of you, there is some uncertainty about statutory reserving. GMIB waiting periods in general have been pushed out fairly significantly. Fund restrictions and a few other things are seen as ways to change the new business side. There are some other things that are going on in the industry. This isn't necessarily about financial risk management, but obviously there's the new tax bill and changes to the competitive posturing of annuities.

Will the consumers pay higher fees? That's the big issue for new business more than anything else. Many of the products and the features that were big sales in 2002 have been pulled. Again, I'd point out that dollar-for-dollar partial withdrawals are probably one of the main ones. GMWBs have obviously gotten a lot of attention and a lot of sales for companies that are on the front wave of that type of benefit introduction. The types of things that companies have been doing to deal with, at least on the new business side, the risks that they find themselves in are discontinuing benefits, increasing fees, tweaking benefits or taking a new approach.

Product balancing sounds good in theory, but it's extremely difficult to quantify. I think there are a lot of other issues around, even if you could quantify potentially the economic impact of having two benefits offset each other or two different product lines offset each other. For example, obviously on the GMDB you have puts in your contract. If you had an EIA block, you'd have a bunch of calls in those products. Some of those exposures cancel each other out, but not all of them. But just purely on the VA side, potential combinations would be, again, offering an enhanced earnings benefit with a GMDB because there you have a call in

combination with a put, or offering a GMIB in combination with a GMDB. You can get more fees from the customer, and they can only do one of the two.

The next one is balancing the VA put options with the EIA call options. In theory, that's one of the types of things that you could view in combination, but, again, it's good in theory but difficult to quantify. You'd have to understand, for example, that maybe it's a good economic hedge offset, but you'd also have to make sure you understood—if you're a stock company or you issue GAAP financials—your GAAP income impact. If one product line is doing its own thing on a GAAP basis, and another product is doing a different thing on a GAAP basis, add those two up and you may not get the product-balancing impact that you would be hoping for based upon the pure cash-flow economic analysis.

On the capital side, the C-3 Phase II capital framework would not allow this specific product balancing—for example, a VA with an EIA. But I believe that some revisions have been proposed to include that possibility in the C-3 Phase II. Once C-3 Phase II becomes effective, it's a little unclear how the rating agencies will react to that. If it comes to a discussion with rating agencies about a product-balancing concept, my experience is that that's a very tough sell.

One possibility of managing risk—although it's not managing the policyholder behavior-type of cash-flow risks, but maybe more capital and reserving or the balance sheet management for a company—is reinsurance. For inforce reinsurance, I'm probably overstating it by saying it's very limited. There was at least one company earlier this year that, in one of their press releases, mentioned they had done a reinsurance of their inforce on a yearly renewable term (YRT) basis. I'll go through an example of where this could help, but, frankly, the main thing it will do is reduce your reserves and capital. When I say that it may increase capital, I'm not talking about the target surplus; I'm talking about your capital on your balance sheet if you can reduce reserves. The disadvantage is that, depending on how it's structured, the reinsurer may not lock in your costs because the reinsurer can't lock in its costs. I'll go through an example.

Another possibility is doing your own reinsurance. What would be the reasons for this? But let me go through some of the other questions. You still need to do some reserves on some type of a basis, probably GAAP. If you're reinsuring through a third party to an offshore that you own, you'd still have to address the capital requirements with the rating agencies, assuming now that that offshore company is not regulated by the NAIC. It may only be a short-term fix to managing your reserves and capital position. It'd be somewhat complex, though not necessarily that difficult. Obviously you're not influencing the underlying cash flows from the policyholder, and you would have to establish letters of credit (LOCs).

The potential advantage of that is that that would be driven by your credit and your capacity because it looks like the reinsurers may be having capacity constraints down the road. Many reinsurers have done a lot of XXX coinsurance and taken

those XXX reserves offshore. They have to set up LOCs for that. If a reinsurer has done a lot of that, they have a tremendous buildup in LOC requirements down the road. They may not have much capacity today to add on stuff that requires LOCs. The advantage of doing your own offshore is that that's driven by your credit and your capacity, not that of a third party who might be already experiencing some capacity crunch there for higher LOC costs.

I'll mention a very brief example to make sure that you get some idea of the implications. I'm assuming now that the charge is just a YRT rate, including, obviously, if a reinsurer is involved, a risk margin component to that. Plus the reinsurer would have to pass its LOC costs on through its YRT charge. If the reinsurer is not willing to lock in its YRT rates for the next 10 or20 years through a reinsurance mechanism with you, the net would be an opportunity where, if its LOCs costs went up down the road, your YRT premiums would go up as well. LOC costs have been going up fairly significantly over the past year or so. Viewed from the reinsurance perspective, the reinsurer might be very reluctant in a reinsurance agreement to, in effect, lock in what it can charge you for the LOC costs.

Here is a very generic example. Say you have stat reserves today of $100, GAAP reserves of $65, LOC costs of 75 bps, and the YRT margin is 20 percent. The expected deaths amount is $4. The reinsurance charge would be $4.80, plus 75 basis points of the $35 difference between stat and GAAP reserves because that's what they have to get an LOC for. So, the LOC charge is, in effect, 26 cents. It's a total of $5.06. The incremental cost to the direct company is only $1.06, because obviously you're going to have those deaths whether or not you have reinsurance. The tradeoff is that it's going to cost me $1.06, looking at one year, but I can free up $35 in capital (although technically that's on a pre-tax basis). That's the type of tradeoff you would have to evaluate. Is the potential for that amount of freed-up capital enough to warrant the incremental cost? Actually, the $1.06 was a pre-tax number, too.

Let me just mention some of the GAAP developments going on in the industry. Again, there are other sessions that cover this in more detail. Here, too, I'm hoping that most of you are aware that there's a new GAAP Statement of Position (SOP) effective first quarter of next year, that will require accruing for investment guarantees that you have on your VAs, whether it's a death benefit, income benefit or any other type of living benefit. It would require some multi-scenario valuation. The actual SOP doesn't go into a whole lot of detail about what that means, but the main change from prior GAAP SOPs is that it's no longer a best-estimate equity growth assumption. You have to do multi-scenario valuation.

Once this becomes effective—I know some companies have actually adopted this—you have to unlock just like you do now for DAC. I'm pretty sure that any changes that would run through as you adopt it would be like any other GAAP treatment. It would be cumulative change in accounting, but then from the first quarter on, any

changes to those reserves and liabilities would be flowing through your income statement and would do any unlocking, again, just like a DAC.

MR. ROGER FROST: When you're modeling GMDB behavior and dollar-for-dollar, logically it seems like someone who has a $5,000 net amount at risk, even though it's 33 percent ITM, would have a different behavior than someone who has a $250,000 net amount at risk. Do you reflect that at all in your modeling?

MR. O'CONNOR: In this modeling we didn't, but that's the type of thing a company would have to look at.

MR. FROST: In our case we did see that we had a lot of exposure at the smaller levels, and it would make sense that we wouldn't have the amount of potential transfers that other people would. Does that seem reasonable?

MR. O'CONNOR: I think so. I'm curious about what companies are doing in terms of monitoring that exposure. For example, I know some companies monitor the activity of people taking out the account value down to some minimum of say $1,000 or $5,000, and looking at that to see if that is really happening or not. If you have a sizable exposure, are your companies monitoring that? At what level are you monitoring it? For example, are you monitoring just in aggregate or do you have the ability to monitor at the distribution level? Do you monitor at the agent level? Traditionally, a lot of policyholder behavior is influenced by the agent.

MR. FROST: We could conceivably do agent level, but we haven't done it. We have seen that most of our exposure is under $10,000 net amount at risk, and we just don't think that people are going to jump through the hoops to protect that like they would for $100,000.

MS. DONNA CLAIRE: We've looked at several companies that are in the market. One of the charts that we did like was exactly what he was talking about. In fact, it's not only down to the agent level, because they were tracking whether or not certain agents were pushing it, but it was basically three-dimensional where it had the amount of "in the moneyness" and the size of it, because it does look like the size probably will have an impact. Size did have an impact in this particular case as to whether or not there would be partial withdrawals.

MR. O'CONNOR: It seems to make sense that, with the larger-size policies, the customer would be more sophisticated, and, frankly, if the agent is going to take the time to go out and try to influence that behavior, the agent will try to do that on the larger-size policies.

MR. GEORGE CHRISTOPHER: I'm going to discuss a case study about dynamic hedging of GMDB risk. The idea is to present one general approach. There are other approaches that you could use. This approach uses dynamic hedging using exchange-traded futures contracts. You need the liquidity of exchange-traded

contracts because you're going to be rebalancing frequently. The first step is to look at a fund-modeling technique. The policyholders have the opportunity to invest in a wide range of funds. You need to be able to connect the risk that they have in those funds with the risks that are out there in the market that can be traded. That's the first step in the approach.

Secondly, we need to talk about scenario generation. It is a stochastic process. You're going to need to look at your parameters and figure out how to fit parameters for a stochastic model. We'll mention that for different uses of the model you may need to use different sets of parameters. Then we'll touch briefly on a valuation model. Performance attribution and the projection system go hand-in-hand. If you can't figure out where you're losing money and why you're losing money or where you're gaining money and why you're gaining money, the confidence level in your hedging program is going to go down. The performance attribution system and the projection system together can enable you to figure out where you'll lose money, where you'll make money and why. Finally, we'll talk briefly on accounting issues, some other benefit types and other products and how the differences between products and the natures of the guarantees might impact how you would go about hedging them.

For purposes of this talk, we're looking at a return-of-premium benefit; your death benefit is equal to the larger of the amount of money that you put in initially and the current account value. A policyholder can invest in a wide range of funds. Hopefully you've been smart about what funds you'll allow them to invest in, and you have a lot of fairly tame funds and not too many funds that we refer to as the "global volatility" fund, where the policyholders are rolling the dice on what's going to happen with their money.

The goal is to take a wildly fluctuating liability that could vary widely from time to time—big claims one month, no claims a few months later—and turn it into a basis-points-of-account-value charge. You're not going to be able to do it completely because the markets are fluctuating. The more that you can make this wildly fluctuating liability look like something that's basically a basis-points-of-account-value charge, the happier people are going to be, because actuaries and regulators understand basis-points-of-account-value charges. These have been around forever. People are comfortable with this. When I say "basis points of account value," it's not an explicit charge to the account value. It's not necessarily connected to any sort of fees that you're collecting. It's used to set up a hedging profit center. The idea would be that you'd fund that hedging profit center with your basis points of account value. They could come from any source. They could come directly from the policyholder. They could come from some other company account. On average, you would break even with your hedging program and you would dampen the volatility quite a bit.

The hedging profit center would execute the trades, track the performance, run the nightly valuations and invest any profits that it might get from the hedging

activities. Those profits might be needed down the road to pay off death benefits. You generate your exposure today, but you don't actually pay any money until the policyholders die.

The first step in the process is to map the funds to hedgeable instruments. The instruments that we're going to look at here are: Standard & Poor's (S&P) 500 futures, one of the more liquid markets out there; Russell 2000 futures, also a fairly liquid market; and NASDAQ 100 futures. You could also use Treasuries, but we're not going to be using them in this example. There are certainly other instruments out there that you could use as hedging instruments. These tend to work fairly well for most blocks of business that we've seen so far. There are a few types of funds for example real-estate-investment trusts (REITs) that won't track the performance of these indices well. We'll talk more about that later on.

You want to develop a set of scenarios reflecting market conditions. There are different sets of parameters that you might need. If you're doing strictly pricing, and all you're doing is looking at what your liability is today and how you are going to go out into the market and hedge that, you'll need something that's fairly close to a risk-neutral pricing model. The reason for that is if the market is pricing using these models, you're going to need to price your liabilities using those same models so that you can get a good offset. If you're not using the same approaches that the market is using, the values of your liabilities and the values of your assets won't move in tandem.

Second of all, you might need some scenarios that are consistent with your GAAP assumptions. It's unlikely that those would be exactly the same that the market is using for pricing. Thirdly, you might need to get a separate set of scenarios that's going to pass the RBC Phase II calibration tests and be acceptable for calculating capital. The same model can be used for all of these. One of the nice things is that since it looks like you're going to be forced to build something like this for your Phase II RBC, you can use it to develop a hedging program also. So, hopefully you can leverage that regulatory requirement into something that'll make you a little bit more profitable or at least make your profits a little bit less volatile.

Finally, you're going to need to develop something for analyzing the performance. There are a lot of moving pieces here. You're going to need to be able to track down which ones of those pieces you're doing a good job at estimating, and which ones of those pieces you need to spend a little more time and effort in going through and modeling.

Why dynamic hedging? There's certainly pressure to do something to manage equity risk. There's little or no reinsurance market out there right now. There are some companies out there who may be entering into a reinsurance market. The ones that we've been talking to would be running the program for you and charging you a few basis points for running it. They'd be doing the same thing that you would be doing in-house, and you'd be paying them a small number of basis points

every year to do it. If you could do it in-house, and you could develop the expertise and develop your own models, then you could save those few basis points of account value.

Equity market risk is not diversifiable. You can't go out and write more policies and get rid of this risk. If you write more policies, it just becomes larger. It's not like mortality risk and some of the other risks that actuaries have traditionally handled where, if you have 1 million policies, your deaths are going to be more predictable than if you only have 1,000 policies.

Static hedges are expensive. Wall Street does not want to write long-dated put options.

Funds are mapped funds to widely followed indices, for instance 50 percent S&P and 50 percent NASDAQ. All funds should have a non-negative amount mapped to each index. In addition, the sum of the allocations to the indices should be 100 percent.

Then you roll the exposures up at the policy level to get index allocations for a policy. You can also summarize the exposure to the indices for the block. If you have a particular index—the S&P 500 is typically one—that there's large exposure to, that might be a good one to go out and hedge. If you have a smaller index, maybe you don't have a whole lot of international funds. So if you use Europe, Asia and Far East (EAFE) as one of your modeling indices, you might not need to hedge that because it's a small proportion of your inforce that maps to that. You don't need to worry about it. Secondly, there is no EAFE futures contract, so, you'd have to hedge that using some other futures contract such as an S&P 500. If you're mapping a lot of your inforce dollars into an EAFE fund, you might want to go out and figure out what else you could map it to that may be a little bit easier to hedge directly.

Ideally, you should do this before you allow people to invest in an individual fund. You should know what the characteristics of that fund are. What indices does it map to? How's it going to track going forward? If the NASDAQ is going up and the S&P is staying right where it is, how are your funds going to behave? If the S&P goes down and the NASDAQ goes down twice as much, which one are you going to be on the hook for? The goal here is to keep the funds out that you can't hedge; keep the funds out that are too volatile. Even if you are hedging highly volatile funds, you're still going to see some residual volatility in the returns.

When you're doing the mapping, you're going to use monthly fund returns. You need the monthly fund total (including reinvested dividends) returns net of investment fees. We fit seven parameters—the six index allocations, plus there's an additional alpha component. Alpha is how much you outperform or underperform your particular allocation of indices. When you're fitting seven parameters, if you only have a year of fund history, you're not going to get a very reliable model. The

idea is that you want to keep funds in there that have some sort of reliable history. Again, the goal in your fund modeling is to make things as stable and predictable as possible. Keep out the funds that you can't map and that you can't figure out exactly how they're behaving relative to indices.

Chart 1 shows a sample of six funds. Some of them are good funds to put in; some of them maybe are not such good funds. Let's look at the S&P 500 Index Fund. We run our mapping, and it says 99 percent S&P 500 and 1 percent T-bill. In S&P Index Funds, the stocks don't get traded frequently, so they do need to keep some cash around. One percent may be a little high. I put it in there just to illustrate the point that it doesn't necessarily map 100 percent to the index because when redemptions come in, they're going to need to have some cash on hand to give out. It's a lot easier to send out some of your cash that you have in a cash reserve than it is to sell the stocks and send that out. This is a good fund to have in there.

If you look at the explanatory power, it has "R-squared" in parentheses. Since we do have our constraints that you have to allocate 100 percent, and you're not allowed negative allocations, it's not a true R-squared. Some crazy funds will give you some non-intuitive results in the explanatory power if you think of it in terms of R-squared. But if you get reasonably behaved funds—anything that you can map to the indices, the explanatory power is going to behave much like a traditional R-squared value. You'd like to see a high value there, meaning that the volatility and the movements in the fund are reflected well by our allocation to the indices.

Then you can compare the actual volatility, which is what you actually saw on a monthly basis in your funds, to the predicted volatility, which is what you would have seen going back historically over the same period, if you took that exact weighting of the indices and then looked at monthly returns and looked at the volatility there. If those two numbers are close, that's a good thing. You can't just look at the explanatory power. You also need to look at those two pieces. Again, you might think that, because of R-squared, you'd only need to look at one, but we have some constraints in there, and so you need to make sure that both of those two things are working together for you here.

Funds 2 through 4 are actual funds out there whose names have been changed. They include a typical growth fund, a government bond fund, and what I called a "balanced fund," which is probably slightly misnamed. A typical balanced fund will usually end up 60/40 equity and bonds, and this one's a little more in the equities. But these funds look pretty good. They all map explanatory powers of 84 percent or higher. If you look at your actual versus predicted volatility, they're all fairly close. They all map in ways that look fairly reasonable. Your government bond fund is mapping to bonds and cash. We don't have anything in there that has some sort of credit spread, so that's not an unreasonable allocation. It's tracking interest rates, but you're going to have to go to the longer-term funds. You have to go further out on the curve, and then the spread gets balanced by mapping some to the 30-day T-bill. It's not a perfect mapping in terms of interest rate risk, but it's pretty good as far as mapping funds.

You might look at the U.S. Growth Fund and see that it has 10 percent in EAFE. Well, U.S. companies do business overseas. You can sometimes see that flowing through in terms of EAFE exposure for a large U.S. company when it goes into a large fund like this. Those first four are good funds; you might want to have those in there. The next one is an REIT. We don't have an index that should map that one well. None of these indices are real estate indices. As you might expect, we get a very low explanatory power, and there's a large difference between actual and predicted volatility. If you're not seeing a big demand for this fund, it might be one that you'd want to exclude because you're not going to be able to hedge it easily. You could also argue, however, that there's not a big demand for it so you're not going to worry about it; it's not going to generate a whole lot of exposure.

Finally, we have a technology fund. It maps 100 percent to the NASDAQ. This is just an example of a sector fund. Sector funds are very tricky because you can't map them well to indices, and you can't really hedge them effectively. You don't get the diversification of risk because you're intentionally keeping all of your eggs in one industry basket or a small range of industry baskets, and that's going to really increase your volatility. That's going to give you exposure on your guarantees, and the policyholders probably shouldn't be allowed to invest in those sorts of funds if you're trying to minimize risk. You're stuck with what's in there for existing products, but these are the sorts of things you might want to try to avoid with your future products.

How do you test your fund model? There are a few basic sanity checks. In our fund modeling sample, if the S&P 500 had mapped to the NASDAQ, that would set off a lot of alarms that maybe this fund mapping is not the best procedure in the world. If your high-quality bond funds don't map to fixed income indices, you have a problem there. If you start looking at junk bond funds, those will start mapping to equity because you're just betting on the fortunes of the company when you start getting into lower-quality bonds. The mapping needs to fit well over different time periods to be effective. Fit it over the last three years. Fit it over four years ago to

one year ago. Look at different time periods and see how stable it is. Stability is the key here.

You want to set up a hedge that you're fairly confident is going to be reflective of how that fund is going to behave going forward. Test it in up and down markets, and in high and low volatility markets. Validate it with out-of-sample data. If you have five years' worth of data, model it for the first four years, and test it on the last one. The model should be stable, which doesn't mean that you automatically have to rule out something because of instability. If you can come up with a good explanation for why that fund is unstable or why you saw a small period of instability, that's probably okay. For example, maybe a new fund manager came in. The first fund manager was underperforming and the new fund manager came in, and that slightly changed the mix and made it a little unstable. It might be okay to put in that fund. Say the fund had a particularly large investment in one of the companies like Enron or WorldCom. That would certainly be an alarm as far as selection, but it would cause instability. If you can figure out why there might be some instability in there, you might be able to include a fund that has a little less stable mapping.

How can you use the results? First of all, you can determine which hedging instruments will be useful. If you select your six indices, and you don't get good explanatory power, then you're probably going to need to change some of those indices. If you've gone through and your customers had been investing a lot in European funds, and you didn't have EAFE as one of the indices, you might get a lower explanatory power. You can go back and try mapping this with an additional index. You can also use it to manage your risks by setting higher charges to the policyholders for different funds. If you have an index fund like the S&P 500 Index Fund, that's going to be quite easy for you to hedge. There are actual instruments out there that will track that same index. You might charge them a lower premium for this.

You start increasing the premium for higher volatility funds. You might want to increase the premium again if it's something that's hard to hedge. So, if there's something that you could hedge, but there's going to be a little more basis risk, you might want to charge them a little higher premium. You might want to exclude those risky funds from the GMDB. These risk management suggestions are ranging from easier, less restrictive restrictions, in terms of keeping your risk down, to more restrictive. Charging a higher fee allows them the flexibility to invest. You're starting to rule funds out as you get down to the bottom. Then the strictest would be to tell them that they can't invest in this riskiest fund, and if they do, then they can't have the GMDB, even if they have only 10 percent in this fund. Again, if your funds seem to be having inconsistent objectives across time, you want to try to keep those out.

Let's move on to scenario generation. You're going to have to select parameters for expected return and volatility. You also need some correlations between the

indices. Multiple scenario generators may be needed. You may need a regime-switching model. You may need also a risk-neutral model and some sort of GAAP model that's consistent with what you did in GAAP. Model parameters should be stable over time. What you don't want to do is set your volatility according to the VIX Index. The VIX is an index that they put out that's based on very short-term volatility fluctuations.

If you're changing the parameters too frequently, you're going to see gains and losses coming through your model just from changing the parameters, and it's going to obscure the useful information that you would be able to get out. Maybe once or twice a year you can look and see if your volatility is still consistent with what's out there in the market. Maybe that's a good time to change those parameters. But if you do it much more frequently than that, you're going to see too much noise coming through just from changing the parameters.

You might need to pass a calibration test, in particular for the RBC. You should set parameters fairly conservatively to avoid losses due to aggressive parameter setting. If you take your model to senior management or to someone outside of the company, and you see an occasional small loss dribbling off, then no one's going to be too upset as long as the swings are fairly muted. But if you've set your parameters, in particular your volatility parameter, too aggressively, you might suddenly see a big spike and a big loss, at which point people are going to become concerned about whether your hedging program is really going to do what it says.

Here's a hedging valuation model. We need to go out into the market and hedge. We'll need to build a trading grid that's going to tell you which instruments you should trade to offset your current risk. Currently we set up a grid that ranges from a 20 percent shock down to a 20 percent shock up. At each 1 percent interval, we will tell you which instruments you would need to hold and which amounts to offset the risk of each index. So we would have six columns, one for each index. In each column we'd calculate your sensitivity to an index at the beginning of the day. The 0 percent line tells you, if there's no shock to the index, how many contracts you would need to hedge the market exposure to that index. Suppose that the market goes up 5 percent immediately. You'd need to readjust your position. Also stored in that table would be the new position that you would have after the 5 percent shock.

Since the liability is not liquid and we're not actually valuing the individual options (we're interested in hedging the block together, not hedging individual policies), we can actually get away with a few things that investment banks might cringe at when you tell them. We can use fewer scenarios. Typically with some sort of a ratchet-type option, you would want to use probably closer to 10,000 than 1,000, if that was the only thing that you had and you had to price that to trade it in the market. However, when we have, for example, 50,000 65-year-olds with the same benefit, we can use 200 unique scenarios per policy. Since we're using different scenarios for each policy, that gets us 10 million scenarios for the 65-year-olds.

That's certainly enough scenarios. You do need to use a few more total scenarios, but it runs more quickly because each policy only has to run through 200 scenarios.

Let's talk about complexity versus simplicity. Our models are currently using just a single risk-free rate and a single volatility number. The real reason for that is that if you go through and you have a full volatility surface like Wall Street says you should have, and a full risk-free curve over time as Wall Street says you should have, you're going to start to see gains and losses from basically riding up or down the curve, which everyone here, being actuaries, is probably familiar with. You'll also see gains and losses from moving around on the volatility surface. It's just noise that would be showing up. Sure, you would see those if you were actually going out and trading both sides of the balance sheet. You don't need that level of detail when you're doing things at the block level; you just need to be close enough that you can dampen the volatility. Those differences from the theoretical value would show up in the hedging program. It's not a big factor there. Again, simplicity is a premium. You need to be able to get these things to run quickly.

How would you examine hedging your losses? Chart 2 shows a fairly complicated equation, but it's very similar to rolling a reserve forward. If you take your option value at the beginning of the period, roll it forward with interest, adjust that for your claims and your premiums and adjust that for any hedging gains and losses from marking your hedge instrument to market, that's going to be equal to your option value at the end plus your gain or loss from the program. It's a fairly familiar process. It's fairly similar to what you do with a traditional actuarial valuation.

For performance attribution: we need to be able to measure the hedging program effectiveness, and we need to be able to figure out who is responsible for what. If your volatility is off, you're executing trades poorly or you're trading too frequently, the hedging center should be held responsible for that. If you have excess mortality or excess lapses, that's probably not something that the hedging trade center has gone through and set up. Somebody else should be held responsible for that.

A good attribution analysis would be able to isolate all of these different sources of gains and losses. If the market is more volatile than you were expecting, you're going to be trading more frequently, and trading hedging instruments is going to generate you losses while you're trading them. You'd like to keep that to a minimum. High volatility forces you to trade, and you get losses there.

If you take the gains from the program, reinvest them and you can't meet your target investment rate, then you're going to see losses there. If you exceed it, you'll see gains.

There's going to be some basis mismatch when you map the funds to the indices, unless you only have index funds in there, in which case you'd have a much less risky product, but it would certainly be much less attractive to your customers.

There's potentially a basis mismatch due to a hedge instrument chosen. For example, if you have some exposure to EAFE Index or some other equity index, if there's not a direct hedging instrument out there, and you decide to model that as being hedged with the S&P 500, those two indices aren't going to hedge as well as an S&P 500 mapped to an S&P 500 hedging instrument would. That will cause you some basis mismatch.

You're probably going to have a few of your funds that are unhedged. If you have a fixed account fund or you have something that's a very short-term bond fund, you're probably not going to go out and trade indices to hedge that. If one of those things starts to underperform, you're going to see gains or losses there.
There are different market closing times. This is a picky one that nobody seems to think about until it happens. The futures market for these indices actually closes 30 minutes before the cash market. You'll occasionally see temporary gains or losses if the market is going down rapidly or going up rapidly at the end of the day. Your futures positions will close, and then half an hour later, after there has been 0.5 percent, maybe a 1 percent move, your cash index will close out. That actually shows up in your performance attribution report. You're going to see a mismatch there. It will correct itself the next morning when the futures market opens. It has to correct itself because when the futures contract expires, the futures price and the spot price will be the same. This is another source of daily gains and losses

Excess lapses in mortality should get a separate piece in the performance attribution. If a policyholder transfers funds from a low volatility fund to a high volatility fund, that's going to impact the valuation and should be tracked separately.

Chart 3 is a projection we did. This is actually done for a ratchet, not a return of premium. The reason is that if you look at a return of premium over the period from 1993 to 2002, it's very boring because people made a lot of money in the early years and didn't actually lose that much. But you can see some of the things that we've done here. For basis mismatch in this projection, we've just assumed that the funds actually do perform as we've allocated the indices. There's no mismatch there. If a fund was 50 percent S&P and 50 percent NASDAQ in its allocation, we're just assuming that it actually does perform that way.

In this projection we were not hedging EAFE and we were not hedging any of the bond indices, and so you'll see gains and losses due to unhedged funds. If you start to see large gains and losses there, you might think that maybe you should start hedging those currently unhedged pieces of your liability. In actual-to-expected interest, it looks like you're earning a little more interest on your invested gains than you were expecting in the model.

As far as volatility gains and losses, if you remember volatility back from the early 1990s, it was a lot lower than it is now. So you're getting some gains there. In the late 1990s, volatility shot through the roof, so you're seeing volatility losses there.

Finally, we have "Other/Unreconciled." In this particular model we don't have any of that. There are two approaches to performance attribution. First you look at how much you lost from basis mismatch, and then you look at how much you lost from your unhedged funds. If you isolate each one of those pieces individually and hold all the other ones constant while you're just looking at that one piece, you might see some unreconciled, and that's due to second order terms. How much did you lose from the fact that volatility moved *and* there was a basis mismatch?

The way we've done it in this example is to layer the changes on one at a time. You'd move through and look at your basis mismatch. That looks great. In the second column, basis mismatch and unhedged funds, how much more did you gain or lose? We're rolling through, sequentially changing from expected to actual. When you sequentially change, by the time you get to the last column, you've actually changed each of the moving pieces from expected to actual. There's nothing left over for the "Other/Unreconciled" bucket.

The projection system allows you to go through and test out various scenarios. It runs very slowly because by the nature of this you need to run very short time steps. You're testing out a hedging strategy. You're probably going to be trading on a weekly basis, so you need to put that in your projection system in order to see the effect of those trades. It allows you to examine the risk that's left after your hedge has been implemented. It also allows you to examine how you would have done over a certain scenario. For example, how would this particular hedge have done from 1993 to 2002? You can use any sort of scenario that you want, such as an increasing scenario or a dropping scenario. The projections can be doubly stochastic. You can say that you would like a scenario to test that is on average, say, up 8 percent but does have some stochastic shocks in there. Why would you want to do this?

Say you're looking at a projection and you're asking: What's the 10^{th} percentile of that projection? You have 8 percent expected but with some volatility around there. What's the 10^{th} percentile? What's the 90^{th} conditional tail expectation (CTE)? You're going to need a doubly stochastic projection. You're going to need one layer of stochastic to handle how you're generating your projection scenarios, and you're going to need a separate level to go through and do the valuations at each point in time. There are some ways that you can get around that, but to be technically correct you'll need to go through and do doubly stochastic. At this point it's just a matter of how many machines you have around to throw at it. If you have enough machines to throw at it, that's great. If you don't, then you'll need to take a few shortcuts to speed up the process.

Your projection parameters are not going to be the same as your valuation parameters. That's not an inconsistency. That's the way that it should be because, again, you could be testing out a particular scenario and allowing to see the effects of increased volatility.

Let's quickly go through accounting issues. Depending on your accounting treatment, putting in the hedging program may actually increase GAAP earnings volatility. Our clients are working on this issue now. It takes some creative thought, and you need to check with your auditors because I'm not an expert on how to handle this problem. You may be able to set this up as an internal reinsurance treaty to get a little better accounting treatment. There's not a lot of guidance on these sorts of treatments because they're all still new. You'd need to talk to your auditor and see what other folks are doing.

Now I'll talk about different types of products and how each would be a little different. As far as EIAs, you start to see more options used as hedging instruments because of the shorter nature of the guarantees here. You can use exchange-traded options or you might be able to get some custom options because it's a shorter guarantee. The investment banks aren't going to be as scared of a one- or two-year guarantee as they are of, say, a 30- or 40-year guarantee. Typically, it'll be cheaper to dynamically hedge in-house, or you could use some combination of the two. If your implied volatilities are low, and you think volatility is going to be going up, you might want to buy options and lock in a low volatility. If volatility is high, and you think it will be coming down, you might want to go with a dynamic strategy so that you can take advantage of the decreasing volatility.

With GMIBs and GMWBs, anti-selection risk is a lot higher. Policyholder behavior is going to be a much larger driver of profitability than it is with a GMDB. I suppose there probably is some anti-selection risk in a death benefit, but you have to die to collect it, which limits the amount of choice going on there. Again, careful product design can minimize these types of risks.

In summary, dynamic hedging can help you. You need to be very careful when you're doing it. You need to have an essential first step with your fund modeling. You need to understand how the funds that people are investing in actually mirror the changes in the market, and you need to figure out what instruments you need to use to hedge those exposures. Several sets of parameters may be needed for a complete analysis. You can't just use one set of parameters for all of your models. You may need projection parameters and valuation parameters. Performance attribution is essential; you need to know which parameters aren't working for you. Finally, different products may require different types of hedging programs, depending on the length of the guarantee and what type of instruments might be out there to hedge your exposure.

FROM THE FLOOR: You said you use 200 unique scenarios for 50,000 cases. Does that mean you actually had 10 million scenarios generated?

MR. CHRISTOPHER: They generate on the fly. It's not like we have a file that has 10,000 scenarios. You tend to key the scenarios off of the policy number, so you can reproduce them. If you send in the policy number, you generate scenarios for

that particular policy from the policy number. For example, on a projection, the system will currently say that we generated a 30-year scenario for it last time. We've moved one year into the projection. We're going to generate that same 30-year scenario, but we're going to cut off the first year so that we are consistent across the policy. You'd be using the same scenarios. I might be using ones that make sets.

FROM THE FLOOR: How do you guarantee that no two of them are going to regenerate the same thing and that it's unique?

MR. CHRISTOPHER: Eventually you loop, but you do get enough variance before you loop that things tend to work out and they're fairly stable. We did a lot of testing on if you change the seeds, and you run things through, how much do the values change? We found that 200 was a good compromise. We get fairly stable numbers with 200, and if we ran more than that, it slowed things down. If you ran fewer than that, your numbers were a little less stable.

FROM THE FLOOR: Does creating these scenarios on the fly slow down the system?

MR. CHRISTOPHER: Not a lot. We have a distributed system, and right now the limit that we're running into is getting data across the network. We use 20 dual-processor machines for our hedging runs, and a quarter of a million policies will run in two to four hours depending on how much network traffic there is. If we were keeping the scenarios around, that would add to the network traffic that we're sending around. That is currently more of a limiting factor than the run-time.

FROM THE FLOOR: I have a question, particularly for GAAP, for purposes of demonstrating the effectiveness of a hedge. In the example we kind of threw away basis risk and a couple of other things. How would you go about that process?

MR. O'CONNOR: My impression is that you can look at a fair-value hedge if the hedged item is on a book basis; it's not already marked-to-market. I think the key question is: Is this GAAP reserve mark-to-market? Therefore, I think that's an accounting issue of whether or not that would meet the requirements of whether or not you could even attempt to get fair-value hedging. That's probably one of the first key accounting issues. Obviously, if you wouldn't qualify for a fair-value hedge or a cash-flow hedge, it's not going to help. Plus, even if you did qualify for a fair-value hedge, any ineffectiveness will still flow through the bottom line. But, arguably, once this GAAP SOP is in effect, the GAAP liability is something that's fair-valued. It has this kind of smoothing mechanism like the DAC, where you look at the retrospective history, but you have a market component to it.

On the modeling, there are other ways that you can do some sampling techniques to try to reduce the number of scenarios run, or on the other side, to increase how quickly they converge. It's a combination of choosing the right number of scenarios

to do and also some sampling techniques that might be able to help you convert things more quickly and run with fewer scenarios.

MR. CHRISTOPHER: There is another thing you could do to illustrate that it works fairly well. We have a parameter that we can put in that adds a random shock between the various funds and the index returns. Then you can go back and look historically. What should that parameter be? How would you perform over five or 10 years if you had added that random shock in there?

MR. ANTHONY FERRARO: When you were doing the fund model and mapping the funds, how did you determine the optimal percentages there?

MR. CHRISTOPHER: We fit the models with a leas-squared-error approach. When you put the constraints in it, it starts to look a little bit different than a straight regression. In fact, if you have a fund that doesn't map particularly well, like a real estate fund, then without the constraints you'd get some really wacky parameters that you would be trying to hedge it with.

Chart 1

Fund Modeling Sample

	S&P 500	Russell 2000	NASDAQ	Salomon Broad Inv. Grade	Morgan Stanley EAFE	30 Day T Bill
S&P 500 Index Fund	99.0%	0.0%	0.0%	0.0%	0.0%	1.0%
US Growth Fund	38.9%	9.9%	40.1%	1.4%	9.7%	0.0%
Government Bond Fund	0.0%	0.0%	0.0%	82.5%	0.0%	17.5%
Balanced Fund	50.4%	16.0%	0.9%	23.0%	3.1%	6.6%
REIT Fund	0.0%	21.6%	0.0%	0.0%	0.0%	78.4%
Technology Fund	0.0%	0.0%	100.0%	0.0%	0.0%	0.0%

	Average Fund Return	Average Absolute Error	Average Error	Explanatory Power (R-Squared)	Actual Volatility	Predicted Volatility
S&P 500 Index Fund	0.9%	0.0%	0.0%	99.9%	15.1%	15.1%
US Growth Fund	0.6%	1.4%	0.0%	89.7%	20.8%	19.6%
Government Bond Fund	0.5%	0.2%	0.0%	86.9%	3.3%	3.1%
Balanced Fund	0.8%	1.0%	0.0%	84.4%	12.6%	11.5%
REIT Fund	0.6%	2.8%	0.0%	16.9%	13.0%	5.4%
Technology Fund	-1.4%	2.2%	0.0%	94.1%	37.0%	28.6%

Slide 8

Chart 2

Hedging Program gain or loss

$$OVt\text{-}1(1 + i) - Ct + Pt + HGL = OVt + NGL$$

Where

- OVt-1 is the option value at time t-1
- i is the earned rate over the period for assets backing the GMDB liability
- Ct is the GMDB death claims over the period
- Pt is the premium allocated to the hedge profit center for the period
- HGL is the gains/losses on the futures portfolio
- OVt is the option value at time t
- NGL is the net gain or loss. This is a balancing item.

Slide 14

Chart 3

Sample Performance Attribution

MG-Hedge™
Financial Projection System
Annual Gain/Loss and Attribution Analysis

$653 Million of Premium issued at end of 1992
Performance of Hedge Program over Historical Return Scenario (1992-2002)

Year	Date	BOP Portfolio Value	Interest earned	Futures Gains	Actual Claims	Actual Premiums	EOP Portfolio Value	Option Value	Net Gain/Loss
0	1992								
1	1993	486,028	22,745	(23,282)	(519)	695,308	1,180,281	503,876	676,404
2	1994	503,876	49,142	4,466	(102,538)	698,127	1,153,073	654,218	498,855
3	1995	654,218	28,222	60,902	(8,771)	756,444	1,491,014	726,878	764,136
4	1996	726,878	53,278	(75,065)	(2,950)	844,235	1,546,376	1,165,521	380,855
5	1997	1,165,521	83,856	(245,282)	(5,427)	946,382	1,945,050	1,555,981	389,070
6	1998	1,555,981	116,722	(277,455)	(21,675)	1,015,970	2,389,543	2,255,404	134,139
7	1999	2,255,404	167,169	(179,899)	(13,725)	1,071,712	3,300,662	3,062,819	237,842
8	2000	3,062,819	232,012	(919,783)	(289,574)	968,213	3,053,688	3,757,420	(703,732)
9	2001	3,757,420	312,080	530,991	(1,567,255)	644,196	3,677,432	4,250,898	(573,466)
10	2002	4,250,898	388,194	2,972,288	(2,352,948)	490,021	5,748,453	6,405,403	(656,949)

Year	Date	Basis Mismatch	Performance of Unhedged Funds	Actual to Expected Interest	Volatility G/L	Actual to Expected Claims	Actual vs Expected Lapses and Transfers	Other / Unreconciled	Net Gain/Loss
0	1992								
1	1993	-	201,600	6,502	468,303	-	-	(0)	676,404
2	1994	-	(329,579)	14,047	814,386	-	-	0	498,855
3	1995	-	285,764	8,067	470,305	-	-	0	764,136
4	1996	-	(81,635)	15,230	447,260	-	-	0	380,855
5	1997	-	48,024	23,970	317,076	-	-	(0)	389,070
6	1998	-	105,816	33,365	(5,042)	-	-	(0)	134,139
7	1999	-	(116,262)	47,785	306,319	-	-	(0)	237,842
8	2000	-	87,784	66,321	(857,837)	-	-	(0)	(703,732)
9	2001	-	(44,851)	89,209	(617,823)	-	-	0	(573,466)
10	2002	-	26,613	110,966	(794,528)	-	-	(0)	(656,949)

Slide 17