2001 Valuation Actuary Symposium
November 29–30, 2001
Lake Buena Vista, Florida

Session 38TS
Complex Liability Modeling Issues

Panelists: Michael L. Beeson
John M. O'Sullivan

Summary: This teaching session covers the development of realistic liability models for use in asset adequacy analysis, asset/liability management, pricing/repricing, financial reporting or other corporate purposes. The primary focus is modeling issues for variable product features, such as minimum living or death benefit guarantees on both variable life and annuities, annuity death benefit enhancements, and variable immediate annuities. Topics include:

- *Analysis of the special risks posed by these products*
- *Selecting realistic dynamic assumptions*
- *Techniques used in modeling these products*
- *Scenario testing and result presentation to best communicate risks*

MR. MICHAEL L. BEESON: I am from Tillinghast-Towers Perrin in Atlanta. My co-presenter is John O'Sullivan. I'll let him introduce himself.

MR. JOHN M. O'SULLIVAN: I'm an actuarial consultant for AON Consulting. I am on the Valuation Actuary Planning Committee and the American Academy of Actuaries Variable Annuities with Guaranteed Living Benefits (VAGLBs) Work Group. If you have any input on any of those two groups, I'm one of the people you can see.

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

MR. BEESON: I've been with Tillinghast for almost 18 years, working with life and annuity products from pricing to financial reporting to GAAP to appraisals to asset/liability management.

We're going to discuss the challenging aspects of modeling products such as variable deferred annuities, variable immediate annuities, and equity-indexed annuities. I'm going to start off with an overview of the various modeling risks, and I'll use examples from variable deferred annuities and equity-indexed annuities. John will then follow with a detailed analysis of the variable immediate annuities.

I'm going to get started by using an example to illustrate some of the risks inherent in these products. Let's imagine a Great Depression scenario. Three months ago or two years ago, it was a lot easier to say it could never happen, but let's assume that we have a 50% loss in funds valued over a 20-year period. Let's further assume, for this example, that our average annual mortality rate is 2%. On average, with a guaranteed minimum death benefit (GMDB) product, the amount at risk would equal the fund value since we've had a 50% loss. If we apply that 2% mortality rate to that risk, we end up getting a charge for the GMDB of 2% of funds. I'm confident saying that no one is selling products with charges equal to 2% of fund. This isn't a very realistic example. That 2% mortality rate comes from attained age 71 or 72 on the male annuity 2000 table, and the 50% drop is not all that realistic. Furthermore, you could argue that even if this scenario happens once in a great while, as long as you have issues spread out over several years, the policies that were just issued before the drop would absorb the brunt of the losses. The policies that had been issued earlier might have had some run up in values; therefore, you're overall risk position isn't quite so desperate.

On the other hand, if you have a ratchet benefit in these products, then every one of those products would have moved up to that high-water fund value right before the drop, so you really don't have that spreading that you were hoping for. You could compound the problem by using a 5% roll-up benefit. Furthermore, instead of using a guaranteed minimum death benefit, where the chances of realization depend on somebody actually dying, you could have used a guaranteed minimum income benefit (GMIB) where the benefit is elective. Consider also a guaranteed minimum accumulation benefit (GMAB), where everybody who is in force at the right time is going to get it without electing anything.

Let me also give you a similar example with respect to equity-indexed annuities. Here your policies have embedded derivatives, and you're probably hedging this with a call option. Because of the money you spend to buy the hedge, you start with less money invested in fixed-income assets. In this Great Depression scenario, your call options are quickly reduced to almost zero in value. Your policyholders lapse and your fixed-income assets haven't accumulated enough to reach your guaranteed minimum benefit, so you're looking at losses on surrender. If this scenario occurs at the same time as an increase in the yield curve, then your fixed income is worth even less than you would have hoped.

The point is not that you have to reserve for a Great Depression, or that you even have to look at the Great Depression as one of your required scenarios in cash-flow testing. What is important is that we recognize the risks inherent in the products, and that our models are capable of adequately measuring or quantifying those risks. We need to be able to convince regulators, rating agencies, a company's board of directors, or even ourselves that we've done an adequate job of modeling these risks and accounting for them.

Our first topic is modeling the past or validation issues. Some products that you might be more familiar with are traditional life or universal life. We're used to modeling or validating such things as beginning fund value, reserves, cash values, policy count, and average size. All this is still applicable when you get to variable products, but there's more information that you need. Take a guaranteed minimum death benefit. What information do you think you need to have in order to adequately model that risk if you have an in-force product? You would want to have at least the sum of the premiums paid, since that's what you're going to be returning. If your guaranteed minimum death benefit is going to be reduced by any partial withdrawals that have been taken, you want to know what that amount is. What else would you need if you added a roll-up benefit? You'd want to know the timing of premium payments. If you're going to accumulate them at 5%, you need to know when those payments were made.

If you then go further and look at an annual ratchet product, you not only need to know the previous information, but you also have to know what your high fund value is to date. How do you get this information? My first observation is that your company has to have this information

in order to pay the benefits. Second, in most cases, you're also going to need this information just to calculate the reserves. It must be available somewhere in the company's accounts, and if you can get it for your modeling, then you're in reasonably good shape. All you have to do is identify the appropriate starting values that you need and go forward.

Unfortunately, we've discovered, in the course of our consulting work, that a number of companies have had problems, at least initially, getting their hands on these data when they're ready to do their cash-flow testing. That leads to the challenging problem of how to recreate the data in a reasonable manner. Basically, what you're talking about is projecting from issue to valuation date using whatever information you can about when premiums might have been paid, how various indexes and economic indicators behaved, and, in some cases, what policyholder decisions might have been.

To give a more extreme example of the same problem, consider the Canadian segregated fund products. It has a return of premium death benefit, as well as a 10-year return of premium guaranteed minimum accumulation benefit. At the policyholder's option, both of these can be reset at any point in time to the current fund value. If you're dealing with this product, and you don't have the information as to the current ratchet value, then you have to somehow try to account for the policyholder's likely decisions between issue and valuation date.

Let's move over to the equity-indexed annuity product, which has similar problems. If you have a point-to-point annuity, you need to know what the index value was at the time the product was issued, so you can determine what benefits are going to be paid at the end of the index term. If you have a high watermark annuity, then you need to know the index value as of the end of each policy year, or at least what the highest value is to this point. For an annual ratchet you have similar concerns.

Suppose you introduce an averaging component to your equity-indexed annuities, so the benefit is dependent on the 12-month average of the index value prior to the crediting date. Next, suppose you're seven months into the final year of the index term. You have five months to go before you're ready to credit this benefit. You're obviously going to model the index for the

future five months, but, as of the model's start date, you're going to need to know either the index value at the end of the previous seven months (so you can do the averaging), or the average value to date. Again, it's much easier to do the modeling if you can get this information out of your valuation system. If not, you're going to have to deal with recreating the data yourself as best you can.

One additional point on the equity-indexed products arises because these are derivative products. Whenever you have a situation where you're not exactly hedged, such as when you bought the options on a different day than policies were issued, your risk is unlimited. The regulators are going to be even more concerned about how well you're modeling your payouts on the liability side and your cash flows on the asset side. You don't have the same limits on the risk that you would have with a simple return of premium benefit.

Another consideration with respect to modeling the past is the granularity of the model. Let's say you have some policies where your guaranteed minimum death benefits are in the money and others where they are not. It's probably a good idea to avoid grouping those policies together because you might obscure the risk. As far as how fine you should make your model, I think most people tend to err on the side of conservatism. Ideally, you should capture all the risk that you need without making your model any bigger than you have to. My practical suggestion is, if you have a model that you think adequately captures the risk, and you're worried about whether it's too big, then save what you've done, try cutting it down and rerunning the smaller version. Then you can determine what difference it makes in terms of the results. If you find that it makes a significant difference, then you needed the granularity that you had in the first place. If not, then you can cut down the size of your model and save yourself some run time.

My final point on validation, modeling the past, is to remember that the real objective is to model the future adequately, but not the past. The future risks are the ones that we're trying to reserve for. I would hope you're not going to spend too much time recreating the past. Again, if you don't have the information necessary to get started this year, and you have to model from issue to valuation date, then go back to your company and try to get that information for next year.

My next topic is economic scenarios. I'm aware that the symposium has devoted entire sessions to the creation of a scenario generator, so I'm not going to give you the two-minute lightning version here. Instead, what I'm going to try to do is identify the important pieces of information you need to get out of your scenario generator in order to adequately model your variable annuities and equity-indexed annuities.

First, let's consider the breadth of scenarios. We know cash-flow testing requires the New York Seven interest rate scenarios, but what about equity scenarios? Is it sufficient to assume a simple 9% growth rate? If you have any appreciation at all for the risks of your products, you'll say no to that question. There is no definitive regulatory guidance as far as what scenarios you should test for equity returns. I've seen various states come up with scenarios that they wanted you to examine for equity-indexed annuities, but even those states don't agree among themselves. You need to consider return and volatility, but not just for generic equity returns. If you're dealing with a product that has several different fund options, whether it's high cap, low cap, international equities, or a bond fund option, you want to make sure that your scenario generator distinguishes the different returns available on each one.

Besides actual volatility, we also need implied volatility in order to price options. If we're going to be buying hedges or selling hedges within the course of our projection, we need to be able to put a price on them, so it's important that our scenario generator be able to provide enough information to do that.

Another consideration is the relationship between interest scenarios and equity scenarios. One approach is to generate 100 interest scenarios and 100 equity scenarios and convolute them to get 10,000 compound scenarios. The problem with that approach is you've ignored any correlation between interest rates and equity rates. Most studies have shown that there's a negative correlation, so as interest rates rise, we expect equity returns will fall more often than not. Why is it important that you reflect that in your model? Remember the equity-indexed annuity example that I gave you before? The worst scenario is when the equity market tanks, your

policyholders are lapsing for the fixed guarantee, and interest rates are rising at the same time. In that case, your fixed-income assets are under water. If rising interest rates and declining equity returns are correlated, we need to make sure that we adequately capture the correlation in our model.

Next we'll look at the impact of future decisions, both policyholder decisions and company decisions. For fixed annuities, we're familiar with policyholder decisions, such as excess lapse and company decisions such as crediting strategy. What kind of different decisions are we going to be thinking about with respect to variable products? First, let's talk about guaranteed minimum death benefits. What decision does the policyholder have to make as far as exercising a guaranteed minimum death benefit? Imagine that the equity market collapses and the guaranteed minimum death benefit is in the money. Other than perhaps seeing stockbrokers jumping out of windows, you're not going to see too much selection against you. For the guaranteed minimum accumulation death benefit, if the policyholder has to be in force to get it, you might see a drag on lapse rates, but I'm not sure there's going to be a whole lot of impact there.

Consider the impact of reset or ratchet options on these benefits. In this situation, you're concerned with the upside. If you have a product that has experienced excellent fund returns the past few years, doesn't have a ratchet benefit, and has no surrender charge, policyholders could decide to surrender the product and buy another one. They've built their own ratchet benefit. You have to at least consider the increased lapses in an upside scenario there.

A case I would like to explore in more detail is the guaranteed minimum income benefit. What kind of policyholder decisions are we likely to see there? Let's construct an example. We have a shadow fund growing at 5%, which is just the accumulation of premiums at 5%. After seven years, the policyholder can elect to buy an annuity with the shadow fund at guaranteed purchase rates. Of course, they can also take their regular fund value and do whatever they want with that, including buy an annuity at current purchase rates. The decisions we are looking at involve the likelihood of annuitization on the downside and the likelihood of lapse on the upside.

I'll state two general principles. First, the more "in the money" the guarantee is, the more likely the policyholder is to exercise it. That would seem to be common sense, but there's still the issue of quantifying it. Second, the presence of commissions is going to have an influence. If an agent or broker gets a commission when the policyholder elects to annuitize, they're more likely to be in touch with the policyholder telling them it's really a good idea to annuitize. This is a good benefit for you because your fund value is worthless now. If the markets collapse, you've got this guarantee, so why wouldn't you exercise it? On the other side, if you don't have the commissioned annuitization, but you have a trail commission during the accumulation stage, you are much less likely to have that broker get back in touch with the policyholder to annuitize and cut off the commissions.

Another question is exactly how much knowledge the policyholders have. Do they have a clue about what kind of immediate annuity they could buy with their current fund value on a current basis, or what their guaranteed annuity is going to be worth if they delay annuitization for a year or two? Did they know when they bought the policy what this benefit was, or did it just sound like a good safety net in case they ever needed it? Having made the decision to buy the product, did they then forget all about the guaranteed minimum income benefit and not even realize what rational behavior is?

Let's continue with an illustrated example. Again, this is contrived and simplified. A 60-year old buys a $10,000 single premium variable annuity with a guaranteed minimum income benefit. To simplify the math, I'm saying there will be no deaths at all until age 82 and then 100% mortality. That way you can punch out all the interest-only annuity factors on your calculators if you want to follow along. The shadow fund is going to grow at 5%, and the policyholder can purchase a guaranteed annuity at 2.5%. Here's our situation. Seven years after issue, the fund value has only grown by a cumulative amount of 10%. At the same time, on the open market, if the policyholder wants to surrender that fund value and buy an immediate annuity, they can get a 5% purchase rate.

Does it make sense for the policyholder to annuitize at this point? Is this option in the money? The shadow fund has accumulated at 5% for seven years to $14,071. The fund value is $11,000. In fact, the answer is that the option is in the money, so the policyholder would get a better annuity by taking the guaranteed benefit than by surrendering and buying an annuity with the cash. Does the policyholder know this? What if the policyholder thinks that the interest rate on the open market is 6% rather than 5%? Suppose we change the example. The fund value, instead of increasing by 10% to $11,000, actually decreases by 40% to $6,000. Policyholders get a much better chance of guessing correctly that their option is in the money and that the annuity is beneficial.

Continue the analysis by considering the timing of the policyholder's decision. We've already shown that the policyholder would do better getting the guaranteed annuity than buying one with the current fund value. Note, however, that this same guaranteed annuity will be worth more next year, because the shadow fund is going to accumulate another 5%, and the expectation of life is going to go down. By waiting, the policyholder can get a better benefit. So just because the policyholder is in the money at age 67, should we assume they will rush out and annuitize? At least part of the decision is based on what the policyholder wanted to do in the first place. My opinion is that we're not likely to see significant premature annuitization. Instead, the policyholder will continue to put money into the fund, and to let it accumulate until planned retirement age. On the other hand, I think you will see an impact on the relationship between annuitization and lapses. Say we're at a point where the policyholder would like to annuitize. If it's a question of surrendering and buying the contract at current rates or annuitizing at the guaranteed rate, if you're in the money, you're going to see more guaranteed annuitizations than surrender and purchase. I'd want to build that into my lapse and annuitization assumptions. Also, if your guaranteed benefit is in the money, but the policyholder is not at the point where they would have annuitized, you could see a drag on lapse rates. The policyholder might realize the benefit is worth something, and it's better to continue this policy than to take the money out.

My generic modeling advice can apply to this situation. If you haven't already encountered this benefit and collected data for it, how can you derive a formula for dynamic guaranteed minimum income benefit elections? Make your best guess, and then test the sensitivity of the profits to the assumption. Consider the process that we used when we were trying to set excess lapse assumptions for fixed deferred annuities. Talk to somebody in your company who understands the product, how it was sold, and what the policyholders are likely to be aware of. Talk to someone who understands the commission structure. Then start rattling off a couple of possibilities. Let's say we're 10% in the money. What do you think the policyholders are going to do? What if the guaranteed minimum income benefit is 40% in the money? After you have received that feedback, and after you've talked to more than one person, imagine yourself trying to set a curve to an exponential assumption. I wouldn't go overboard trying to fit a curve to the data because it's all guesses anyway. Take your information, make a good guess, and then once you've done that, consider the possible alternatives. Find out how sensitive your results are to this function. If it doesn't matter, then use your guess. If it does matter, do more research. Try and find a way to reinsure, hedge, or otherwise manage the risk. Regardless of what you do, if you've gone through the exercise of testing your best guess and testing your sensitivity assumptions, you'll at least be able to make a better case to the regulators and the rating agencies that you have adequately considered the risk. We may not know what the exact formula is going to be, but we at least know how sensitive it is.

A final observation about the guaranteed minimum income benefit is that to the extent that you really don't know what the lapse assumption or annuitization elections are going to be, it makes it a little bit more difficult to buy an appropriate hedge against downside market performance. That's because you don't know the timing of the cash flows and when you're going to need the option to pay off.

Let me switch over to equity-indexed products. We have the same issues. For this product, I am more worried about both policyholder decisions and company decisions. Company decisions are going to come into play whenever you need to reset liability perimeters such as participation rates. Let's take an example. I have a nine-year product, a one-year annual ratchet, underlying guaranteed values equal to 100% of premiums accumulated at 3% per year, no surrender charge,

and a participation rate that starts at 50% per year. You have the option to reset the participation rate each year. You developed a budget of 5% to spend on options to support this product based on how much it costs to cover your underlying guarantee and the initial cost of options. If you try to model changing participation rates, what are you going to consider? What's going to change the cost of the hedge from one year to the next? Things like changes in implied volatility or changes in interest rates could make options more or less expensive, and that should determine what kind of participation rate you're willing to offer. The projected experience as you go through your scenarios will change things like the strike price you need as the relationship between your underlying guarantee and your current fund value will determine how expensive the hedge has to be. Also if fixed interest rates change, you might not have that 5% to support your option budget because you might need to more or less support your underlying guarantee.

How important is it that we model the participation rate reset for cash-flow testing? You might argue that because you have the freedom to reset the participation rate, you only have to model up to the point of reset. If you're going to go beyond cash-flow testing, and you want to show that you're managing the risk properly when you talk to rating agencies or company management, you really should consider what's going to happen when you have to reset a participation rate. Otherwise, if you assume a constant participation rate, regardless of economic conditions, you will introduce unrealistic volatility into your projected results. You'd be introducing a mismatch between your option costs and your benefits.

I want to talk about the interplay between lapses and investment strategy for equity-indexed annuity products. How do we normally set dynamic lapse assumptions for equity-indexed annuity products? Often, we use a base lapse assumption, and then we increase it dynamically if the equity market is underperforming, and maybe reduce it if the market's getting better. I have no problem with that basic approach. What I wanted to look at though, as I said, is the interplay between the lapse assumption and the investment assumption. Sometimes companies will try to estimate the amount of options they need to buy based on their expected persistency. Let's say you have a seven-year point-to-point. You assume 3% lapses every year. At the end of seven

years, only 81% of your policyholders are still around to get their equity-linked benefit. If you rely on this assumption, you may limit your hedge purchases to cover only 81% of the fund value. I'm not going to comment on whether this investment strategy or any other is good or bad. My point is if your cash-flow testing uses the same base lapse assumption of 3% (the exact same assumption that you used to determine your hedge purchases in the first place), you've created a projection that's going to give you the least possible volatility. If you have a projection that says you are assuming 3% lapses, and your hedging strategy is based on 3% lapses, you'll get to the end of the seven years, your hedge will work perfectly, you'll get no volatility, and you'll get a terribly misleading result.

My recommendation is if you've built a lapse assumption into your hedging strategy, then, when you do your cash-flow testing, introduce some kind of expected deviation in your lapse assumption in both directions. Look at your past success in predicting lapse rates, comparing experience to pricing, comparing last year's model to what actually happened, looking at fixed annuities, or whatever you can do to get some idea of how successfully you can forecast for these lapsed rates. Then take your typical difference and throw that into your equity-indexed annuity cash-flow model. If you're usually off by 1%, then run your cash-flow testing model using a base lapse of 4% and a base lapse of 2%. See what that does to your results. If you just go with that base of 3% that matches your investment strategy, you've failed to consider the volatility in your lapse estimate. A sharp regulator is going to recognize the problem and go back and talk to you about it.

This brings me to my final topic—modeling reinsurance. I think the temptation is to say that if we've reinsured the benefit, we don't need to model it; it's off our books. Let me give you a couple of examples. We'll see how well hedged we are. Case one is a variable deferred annuity, guaranteed minimum death benefit. The reinsurer's going to pay the excess of guaranteed minimum death benefit over fund value, and they're going to charge you a yearly renewable term (YRT) rate fixed in the contract times the net amount of risk at the beginning of each month. So are you hedged? The answer is you hedged the mortality risk because you have the fixed YRT

charge. You don't have to worry about whether the deaths are going to be 4% instead of 2%, but you really haven't hedged the market risk. Remember the charge was YRT times net amount at risk, and the net amount at risk varies according to your fund performance. Therefore, the amount you have to pay is determined by how much in the money your guaranteed minimum death is. In this case, we've reinsured part of the risk, but not all of it.

Case two involves equity-indexed annuities. The reinsurer is going to pay all the indexed benefits. If there are excess surrenders during the course of the index term, the reinsurer will pay the market value of the hedge options that have been purchased to support the index-linked benefits that are lapsing. How well are you hedged here? The answer in this case is you've hedged the index-linked risk because the reinsurer is paying that whole thing. What you haven't done is hedged the underlying guarantee. You still have the problem that if excess surrenders occur and the supporting options are worthless because the equity market has tanked, you don't have any coverage if your fixed-income assets are inadequate to support your guarantee. I'm not trying to say that anybody made a mistake with their reinsurance. These are perfectly valid reinsurance choices, and I'm not here to recommend what you should or shouldn't do with reinsurance. The point is that you need to understand what it is you've got, what you've adequately hedged, and what you haven't. You don't just say it's reinsured, so I'm not modeling it. Instead you need to make sure that you are modeling whatever risks you're retaining.

I think I have time for a lame joke to illustrate a point. Let's assume you're shipwrecked on a desert island, and you've managed to salvage four items from the ship: a life preserver, a waterproof deck of cards, the ship's cutlass, and an iron frying pan. Which one of these is actually going to be most beneficial to you in terms of being rescued? The answer, of course, is the deck of cards, because all you have to do is start playing solitaire, and someone will appear behind you saying you could put the red eight on the black nine. The point I'm trying to illustrate is whether you're playing solitaire or performing cash-flow testing, there's always going to be somebody looking over your shoulder. It's important that you not only cover the risks, but also demonstrate that you've adequately modeled the risks inherent in your product.

With that, I think I'll turn the discussion over to John O'Sullivan.

MR. O'SULLIVAN: Usually at these sessions you see an awful lot of discussion about variable deferred annuity and modeling issues. I decided to go ahead and spend part of the time talking about variable immediate annuities or variable income annuities. This is something that isn't typically covered in these sessions. It is also a very opportune time. Many companies have started emphasizing defined-contribution money over defined-benefit money. There is an awful lot of money that's pooled up inside of IRAs. In addition, American workers love equities. We're going to spend a fair amount of time in that area.

I'm going to talk a bit about mortality because it's easy to ignore it and to lose track of the importance of it. I'm also going to talk about innovation and some of the things that are in the marketplace, primarily with respect to income determination and guarantees. I'm going to talk a bit about stochastic modeling and about the gain in death benefits at the end to sort of illustrate one-sided versus two-sided risk.

A generic variable income annuity would typically offer a choice of assumed investment rates or benchmark rates of return. There are a variety of income plans that you could have, such as nonlife contingent or life contingent plans. The nonlife contingent would be something like a 30-year variable annuity certain. It is not different than a series of partial withdrawals from a deferred product, and there is typically very little mortality risk. As for life contingent plans, you could have it covering one life or two lives. Obviously, that affects the profitability, as does what you're paying upon the first death in the case of a joint life, whether or not it's a reduced payment or it continues at the same original level.

The other thing that you should always look at is how the monthly income or the periodic income is determined. That's an area where a fair amount of innovation has happened, and there are a lot of differences between companies. You also see liquidity features and whether there is any kind of secondary guarantee on the income payment.

When dealing with mortality in a variable immediate annuity, you're really playing with three different sets of assumptions. You have the valuation basis, which is really where your statutory reserves are going to be based. You have what's in the pricing for the customer or what he was charged. Finally, you have the experience basis. You have to think more in terms of life insurance and less in terms of deferred annuities. I think we tend to be a little bit sloppy about the mortality assumptions. In the income annuities, one does have to go ahead and look at not only where the mortality level is now, but also what's going to happen with future mortality improvements. The A2000 table is the current valuation table being used for income annuities. It continues a very bad practice of not incorporating any kind of mortality improvement into the projected mortality rates. That's something that's easy to overlook, but you really shouldn't. It's almost an unknown. Nobody really knows what's going to happen with future mortality trends. Historically, there has been a significant improvement in it.

I'm going to illustrate something using a particular set of facts. I'm assuming that it's a male, age 70, it's a single life payment with a guaranteed period of 15 years that is also life contingent. The income is $100 a month—I don't know what a person's going to do with $100 except buy a couple cartons of cigarettes. There is an AIR of 3.5%. I'm assuming it's issued in 2001. I'm using the A2000 table. What I did for the experience basis was I assumed that it was the basic 2000 table, and I projected for future mortality improvements. The scale that I used was lifted off the group annuity reserve table, which is such a great report to look at if you're going to be doing some work in this area. It gives you a much more robust feel for what happens with annuitant mortality because the group people have a lot more experience than the retail folks do.

Let's look at this 15-year certain with life payments in Table 1. If I look at the present value using the valuation mortality for the net single premium, the payment is $16,622. The income payments for the first 15 years represent 85% of the total premium. You can see where the liquidity feature is going to come into play. If you allow somebody to withdraw the non-life contingent portion of the payment, you can introduce fairly high levels of liquidity in your early years, and you don't have to worry about the mortality antiselection.

TABLE 1
Mortality Element

	Payments	% of Total Premium
LC vs. NLC		
• NSP for All	$16,622.29	
• NSP for 15 Year	$14,041.23	85%
Val vs. Experience		
• NSP Val Basis	$16,622.29	
• NSP Exp Basis	$16,580.39	
• Mortality Gain	$41.91	0.3%

What can I expect as far as the statutory mortality gain or loss? Table 1 shows a slight gain of $42 at 0.3% of the premium. If I dig a little bit deeper, and I take a look at the incidence of my mortality gains and losses, it's really a strange pattern. You get early gains and losses in the middle years. Then you have gains in the later years. What you have as an interplay is 10% loading over experience on the A2000 table. You're pretty close to the base year for the valuation table, so that projection of mortality improvement is not the dominant feature. The load is the dominant feature. When you get into the middle years, the mortality improvement is overcoming the load that is built in. In the later years, you assume that older, advanced ages will cause less mortality improvement because everybody will die at some point. When you see this for the first time, you might think it's a very strange animal, but it does happen.

Going on to the next step. Let's say that's the pattern of mortality gains and losses. I want to look at this liquidity feature. As I mentioned earlier, you might allow people to take out just the non-life contingent payments, since the payments are made one way or the other. That doesn't expose you to any kind of mortality antiselection, but there are three liquidity features that you can see in the marketplace today that expose you to mortality antiselection. The mildest is the reversal of the sale in the first 12 months. The average ticket on these variable immediate annuities is pretty good. It's about $130,000 or $150,000. That's not unusual for an average size contract to be issued. When somebody gives you that amount of money, I don't think they want to hear that they can't get any money back. The idea here is that people are unfamiliar with income annuities. If you have some sort of a provision that will allow people, within the first year, to reverse the transaction, it raises their comfort level a lot.

The second kind of design feature that you'll sometimes see will involve a mortality antiselection piece where you allow people to withdraw the life contingent part of the payment. What the customer has bought is a life-only plan; there's no certain period on it at all. There is at least one product out there that will allow people to come in once a year and withdraw a portion of those payments. All future payments are reduced. You can see that that's going to expose you to some mortality antiselection.

The third design feature takes a little bit of explaining. It's a leveling kind of concept. We talked about the 15 years being the certain period, and the pure life contingent payments came later. If somebody comes in and wants to withdraw money, what happens is that, traditionally, it's the non-life-contingent payments that get withdrawn. Those payments are in the first 15 years. That's sort of like a step function. I'm actually doing it backwards. In the early years, it's down, and then it pops back up. This leveling concept takes some of the money from the back years and fills it into the early years. In all these types of features, the mortality antiselection is not insignificant. The 12-month reversal makes the purchase less scary for customers. They receive a current account value, which is the amount of money they gave you, adjusted for investment performance, adjusted for the income payments you already made. What they should be receiving is the present value of the future income. Determining the mortality weighting in that present value calculation is really where you have to make the judgment call.

Let's do a little arithmetic to sensitize ourselves to the risk. If 90% of the expected deaths in the first-year come in and need to reverse the transaction, then that's going to cost $43.70 in mortality loss. Basically, a mortality gain is wiped out at that point. If I'm still within the first year, but my health has changed, wouldn't it be a smarter thing for me to take my money out. I'm assuming there is some kind of a spillover effect for the deaths that would have occurred in years two and three. I'm assuming that 50% of the people that would have died in the second year know, in the first year, that they are sick. They are going to go ahead and pull out of this deal. Similarly, with 20% of the people that would have died in the third year. You can see that including the mortality antiselection turns a modest gain into a modest loss. Depending on your

definition of modest, it's about half a percentage point of the net single premium, which is probably not modest anymore. The point I want to make here is it's really easy on these things to say it isn't going to cost that much. Until you go ahead and start fooling around with what some of the numbers are, you really don't know.

One of the areas of innovation is really how you determine the amount of income. When I studied this stuff, we had one plus the actual return in the numerator, and in the denominator, we had one plus whatever the benchmark rate or the assumed investment rate is. That's a ratio we come up with. We multiply that ratio times the annuity unit value, and come up with an annuity unit value. We basically go ahead and pay that number of annuity units and my income varies monthly. That's not necessarily the case any more. People like to get paid monthly. None of us would like to go ahead and just have an annual paycheck. Retirees are getting a monthly paycheck as long as they live. Monthly volatility is a lot more than proportional annual volatility. The idea here is to find a way to smooth out this monthly variation in the income payment. There are really three ways that you can do it by broadly classifying it. First, historically, the separate account looks like it's making an annual income payment, but instead of making that annual income payment to the customer, it makes it to the general account. The general account goes ahead and funds a one-year annuity. At the end of the year, it gets a purchase order again with money coming over from a separate account for another one-year type of a payout.

A second kind of an approach is where all the smoothing is done inside of the separate account on the fund. I'll go into a little bit more detail with that in a minute. In the last one, it's incorporated into the design itself. It's almost like universal life mechanics, where you have a fund and the fund basically determines when the income payment is going to be adjusted. If it hits a lower boundary of the range, you're going to reduce the income. If it hits an upper boundary, it's going to increase the income. You can make this stuff very complicated if you want to.

The idea behind the separate accounts smoothing approach is that you start the contract, paying the customer the nominal income that they would, but you pay it for the full year. You take the money out of the separate account to pay them. At the end of the year, you compare the amounts

you have in the separate account with the amount that would be the tabular reserve or the tabular fund. Then you basically true them up by balancing to the new income amount. With the separate account smoothing, in the first year, you can see it's in between the minimum and the maximum amount that would have been paid, so it's an average. The same thing happens with the second year, and the third year.

In Chart 1, the jagged lines going up and down reflect what's happening if you were to determine and pay the benefit monthly. You're paying it monthly in any case. The straight lines are where the payments would come out with the separate account smoothing. There are a lot of nuances in these kinds of products. The smoothing device can affect your overall profitability, so it's not unusual to see slightly higher profitability with the smoothing because you're getting to hold the assets a little bit longer, assuming that you have a positive investment return.

This kind of gets us very naturally into talking about floors. You can have smoothing that is independent of floors, but if you do have smoothing, it does impact the cost of your floors. I'm going to get into that a little bit more. A floor is a percentage of the initial payment. It's basically a flat amount that doesn't move as long as the customer owns the contract. It usually requires a choice of a fairly low benchmarked rate of return or an assumed investment return. In this case, 3.5% is what we typically see. As I mentioned, it's impacted by the smoothing design that you use. There are often fund restrictions so that you have to invest in Standard and Poor's (S&P) type of fund. It might be just a straight indexed or an enhanced-indexed type of fund. There are some products out there with ratchets in them, but the ratchet doesn't work like a guaranteed minimum death benefit ratchet on a variable deferred annuity. You don't give away all of the upside. Basically, your payments would have increased. The payments increase, but the floor increases at a slower rate.

Think about the traditional way of determining the monthly income payment. What you're saying is if the payments go up in month one, that higher amount gets paid to the customer. If, in month two, the payments go below the floor, we're going to make up the shortfall. You can see

that the customer with the traditional method of income determination gets the benefit of the upside, and you can pay the downside. With the separate account smoothing, any difference in shortfall is being spread out over all the future years. It's going to affect the floor costs. Even though you might make payments now because of the floor, you might be able to recoup them in later years by holding back on customer increases. That's good from your risk profile and from the cost of your benefit. The other side of which you must be cognizant is you have a customer disclosure issue on your hands too.

I wanted to model this stuff stochastically to show you the impact of these two different ways of determining income on the floor guarantee. I did the easiest thing I could—a lognormal distribution. I went to the VAGLB report and pulled off the mean and the standard deviation for the equity class side of a report in 2001. I assumed what is a pretty typical basic annuity charge of 125 basis points and a floor guarantee charge of 100 basis points. I did 1,000 runs—first with the traditional determination of income, the floor guarantee at the initial payment amount, and with a 45-year horizon. On these kinds of income products, you need to use a longer pricing horizon with a deferred annuity because you don't have the lapses, which means you have significant amounts of in-force business. In the case of this kind of smoothing device, the costs of the floor guarantee are deferred. To the extent that your money is in a separate account, you're using that money to pay the customer. If you paid out more money, you'll recoup it off of future income payments to the degree that they're above the floor. Both of those things push out the benefit costs to the later years, so I wanted to capture the flavor of those things. I went up to the terminal age of 115 of the table. My results showed that there was absolutely no cost 15% of the time, and the costs were over 1% of the charge 5% of the time. Chart 2 shows that as you go along, the line tends to bump up a little bit, but stays pretty low. Around scenario 950, the line goes over the 1% charge. My methodology here was a reflection of my laziness, in part, and my technological skills. I basically picked a before-tax interest rate that's reasonable. I looked at the yearly cost of the shortfall, which shows that I'd have to make up the income payments, which I would present-value, and then I'd present-value the asset base. I then come up with a basis point equivalent.

Then I went ahead and did it using the smoothing kind of technique that I talked about where you leave it in a separate account (see Chart 3). For close to 90% of the time there is no cost, but once you go ahead and run out of money in the separate account, you're really in trouble. You have only about 2.5% of these scenarios running over the 1% charge. You have a significant reduction in the cost of your floor guarantee when you use this kind of approach. The impact of the floor guarantee design on cost really shouldn't surprise anybody. If you look at what we've been doing in the deferred annuity world, the design of the secondary benefit has a material impact on what the cost is. It's the same principle, and it is just applied to the immediate annuity side.

The next point I wanted to make is to illustrate that the lognormal results aren't all that realistic. Bad things happen more often than one would expect. I took large company U.S. stocks, and I assumed an average mutual fund expense ratio with the same product charges. I used the traditional calculation of income and looked at something that was issued on January 1, 1929. It was not a great time to buy a variable immediate annuity if they were available. Chart 4 shows the monthly income over 25 years. The poor guy never got back up to $100, which is where he started. If you did this business using the separate account smooth or the traditional approach, you'd have some real cost.

The next thing I looked at was a purchase 40 years later (Chart 5). This captures the 1973 and 1975 bear market. This was not a great time to buy a variable immediate annuity either. I mean if you didn't have a floor guarantee, you were essentially looking at (with the 3.5% AIR) a period of 15 years before your payments got back to where they originally had started. Under this type of an investment scenario, the smoothing thing is going to have less of a cost because you've had favorable investment performance beyond the 15 years. You can recoup some of the money that you paid out already.

One of the points I wanted to make is that the independent lognormal doesn't work well at the tails. This might or might not affect your reserve. It depends upon whether you think the 83-1/3 is out on the tail. If you're looking at risk-based capital (RBC), or the way you're pricing it is

going to make a difference. Take the VAGLB report for June 2001. We have the first derivatives matched, but we don't have the third and fourth derivatives matched. That's the problem. If we use the lognormal, the skew and the kurtosis are neutral. That is what's causing our problem.

I'm going to kind of switch gears and grade into some of the stuff that we've done on the VAGLB committee. First, the proposed guideline does affect these guarantee payment annuity floors, or the GPAFs as they are called in the report. We started by saying that if you had to do something other than the Keel Method, you'd use the lognormal. We gave a mean, and we gave a standard deviation. I think the realization hit us that this probably was okay at the 83-1/3 percentile. There was work happening on the RBC committee for these things, so more sophistication was needed for it. Rather than specifying a particular distribution with the mean and a standard deviation, calibration points were established. Calibration points essentially show what the cumulative investment results or the accumulation factors would be. That's sort of the lead points that you can go ahead and specify.

The quad M report also talks about some principles that an actuary should follow in setting up distributions, which are some other constraints on you. The first, fifth, and tenth durations really need to be met. Since we're working with the 83-1/3 percentile for the reserves, that was a natural one. We also have the 16.7 percentile, and that gets into the question of one-sided versus two-sided.

I'd like to illustrate what's happening by looking at the total equity class and the accumulation value of the dollar. If you would have ranked the returns, the 16.7% shouldn't be more than 0.98 and the 83-1/3 should not be less than 1.3. You have a somewhat normal shape. You're pulling the ends apart to get more into the tails, which is basically the view that makes the most sense to me. Then you get into this question of whether you need to worry about both tails or just one tail. I have examples. For a 5% roll up, the risk is underperformance, and I think about that as a left-tail risk. Any kind of an annual ratchet guaranteed living benefit on deferred annuity has a

different risk. You get great performance early on, and then you get very poor investment results. You have both overperformance and underperformance as your risk. There are some other products out there where the risk is really overperformance and the prime example of that is a percentage of the account value death benefit.

I thought I'd talk a little bit about the gain death benefits because there are some subtle differences among the designs that folks have out there. The general form is that it's a percentage of the gain. The percentage is typically 40%, but the gain can be determined in either of two ways. You can determine the gain (which is relative to the premium) as either what the account value is or what you would pay out to the customer. It would be the account value plus any guaranteed minimum death benefit (GMDB) net amount of risk. It is subject to a maximum.

What I wanted to show is that if you just base it on a percentage of the account value, you really have upside risk. When the market goes down, your risk goes down. If you have it adjusted for the GMDB or net amount of risk, you have both upside and downside risk. This would be the net amount of risk that would be paid out on the benefits. I've got some assumptions there, but that's the point to make.

In these types of situations, we were always asking ourselves how to calculate the statutory reserves. There isn't always a clear answer. The gain death benefit is probably the clearest example. In that particular case, AG34 might not be appropriate. In fact, there's wording in there that says, in effect, you shouldn't use it if your net amount at risk increases with good investment performance. However, there are a number of starting points that have been established. There is AG33 and AG34 for the guaranteed minimum death benefit. Quad M has many more principles. We're in process of revising the practice notes to Quad M, so if anybody wants to see something pinned down, let me know or let one of the other committee members know. We're going to try to fix whatever glaring omissions there are in the near future. We'll probably have something released in the first quarter of 2002.

In all this, there are some principles that are emerging. One of the principles is that it's an integrated benefit reserve type of structure. You look at all of your benefits relative to all of your revenue that's coming in. Another principle is to establish reserves at the 83rd percentile. I'm afraid that in one way or another we're all going to be getting into this stochastic stuff. I think the main constraint is to be practical. Convince yourself that what you're getting is a sufficient look at the risk. There are three methods in the UAGLB Report. There are full stochastic, and the representative scenarios, which are really just a simplified stochastic. It's analogous to full stochastic on cash-flow testing versus the New York seven. Then, try to find some sort of closed form type of thing, like the Keel Method. It is obviously the easiest thing to use. Beware that it might be overly sufficient. You might be costing yourself a lot of earnings because you might be holding too much. You still end up going through a stochastic type of thing.

I want to recap. My advice on the variable immediate annuity is don't ignore any antiselection potential. The design choices that you have can have an impact on related matters, like the floor guarantee. Profitability will vary by income plan and types of liquidity options. You should design your choices to tackle how the income is going to be determined before you get into any of the secondary guarantees. It's obviously important. On the front of stochastic testing, be real careful because it's easy to go ahead and get very precise results, but they might not be correct. There's stuff out there on the CIA website on Segregated Fund Guaranteed fund. The Academy website has been greatly improved over the last year, so I'd urge you to go ahead and look at that stuff. The practice note has a bibliography. If you have a benefit that doesn't fit into one of the actuarial guidelines, there are a lot of principles out there that you can use to pull together and come up with something that's very workable.

MR. JOSEPH M. RAFSON: I'm with Allstate Life. My question is for you Michael. All the issues you pointed out pertain to some of the death benefit and the income benefit designs on variable annuities. Even though these issues are giving us problems in getting the data and the issues, the real problems we're trying to deal with are software limitations and data limitations around the dispersion of both fund returns and fund investments and the dispersion of investment

aggressiveness, where averaging fund returns essentially eliminates your benefits. You might have a shrewd investor and a less fortunate investor. One is way in the money and one is way out of the money. If you average the two, you essentially eliminate your death benefits or your income benefits. How have you seen companies begin to address these issues?

MR. BEESON: I guess all I can say is that it's an issue of granularity of the model. The first question is, can you gather the data yourself to have an idea of what your profile is among your policyholders? If you don't really know your starting point, all you can do is take your best guess as to how you can model it. However, if you can get an idea, then I think what I'd try to do is look into how I could group the results in terms of how much in the money they were at the beginning of the model. I'm taking an analogy from the way we model universal life. We do things like stratify the model by fund value. Here the fund value is very small, and so most of your margins are going to be mortality gains. Here's one that has very high values, so most of your income is going to be in interest margins. I'd be looking here at the extent to which your various policyholders are in the money. This might be a case where, assuming you have the data to start with, you begin with a lot of detail in your model. Then you start trying to selectively combine cells and you justify reducing your model without having any impact on the results.

MR. THOMAS MITCHELL: I want to quote your slide. How would you model volatility, both actual and implied?

MR. BEESON: I intentionally didn't address that because I would rather refer that to my colleagues who are talking about how to design scenario generators. The short answer for me is I get our scenario generator, which is developed by experts who have experience modeling assets and looking at economic interactions and relationships. They'll look at historical volatility, but you have to deal with issues such as the length of the historical period to use to develop the parameters. For example, we've seen volatility increase significantly in recent years. If we look back over twenty years, ten years, or five years, we'll get different answers. I don't have any hard and fast guidelines about how to generate implied volatility yet. I'm hoping I'll be able to learn something from reading the transcripts of some of those other sessions.

CHART 1
Smoothing Income Payments

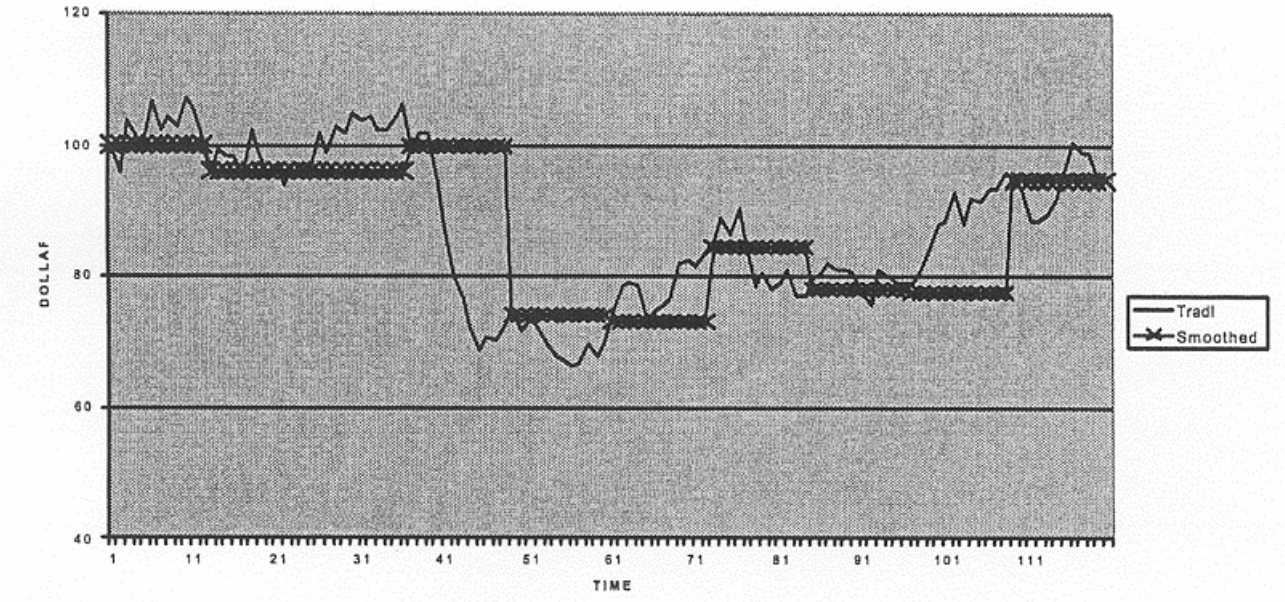

CHART 2
Stochastic—Traditional

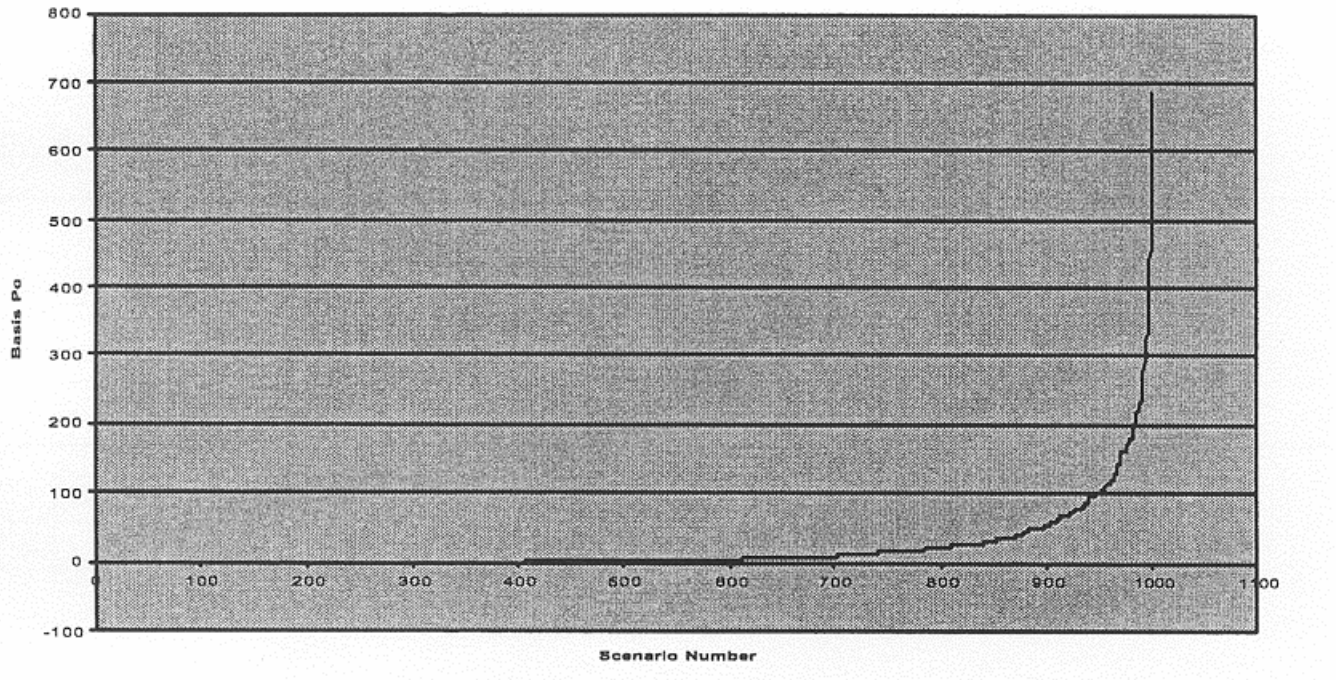

CHART 3
Stochastic—S.A. Smooth

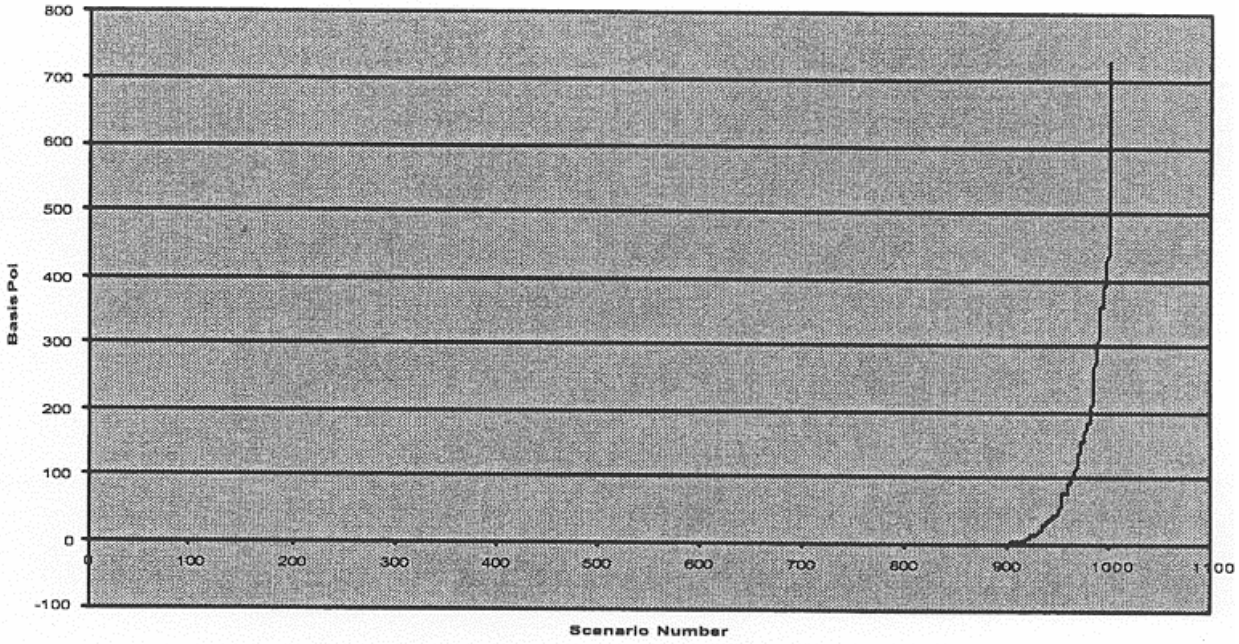

CHART 4
Monthly Income—1929 Issue

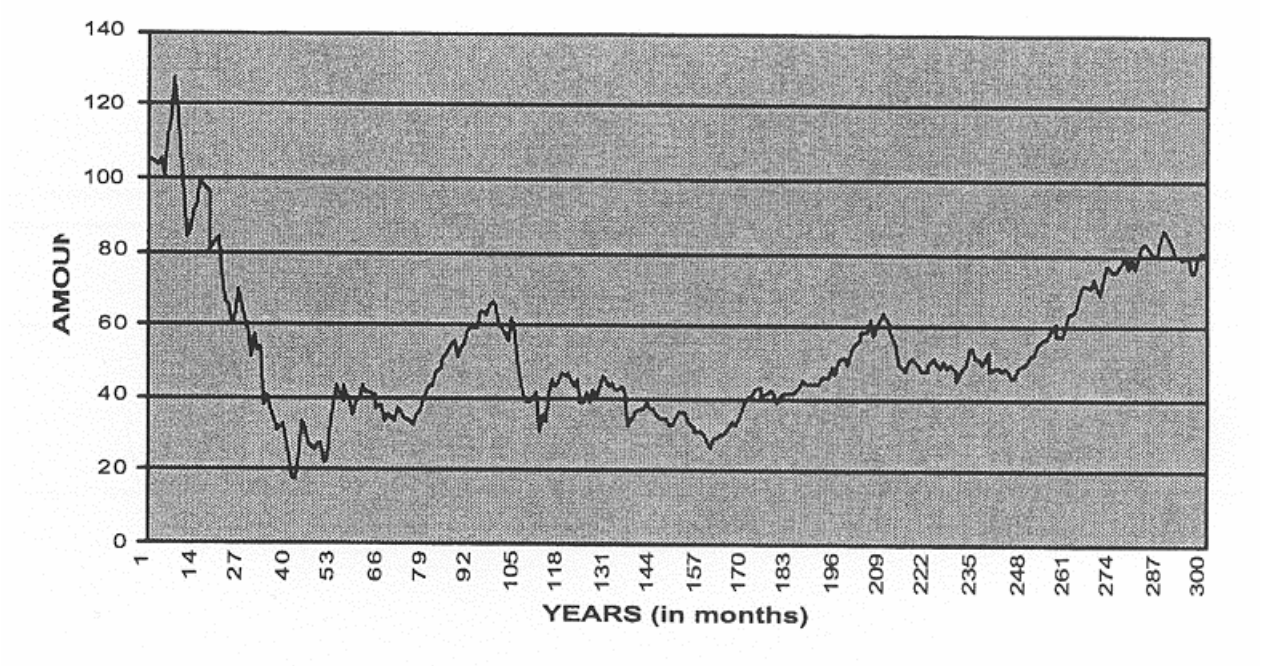

CHART 5
Monthly Income—1969 Issue