2000 Valuation Actuary Symposium
Washington, D.C.
September 14–15, 2000

Session 17PD
Minimum Guaranteed Benefits on Variable Annuities

Moderator: John M. O'Sullivan
Panelists: Franklin C. Clapper, Jr.
                 Charles Dana Tatro

Summary: There has been significant innovative product development in the area of guaranteed minimum death benefits, guaranteed minimum income benefits, and guaranteed minimum accumulation benefits on variable deferred annuities, along with floor guarantees on variable immediate annuities.

This session describes the types of benefits found in the marketplace along with the associated risks. Current reserving practices, risk management techniques, and the latest American Academy of Actuaries (AAA) and National Association of Insurance Commissioners (NAIC) developments are discussed.

Specific examples are provided to familiarize attendees with the "keel" method and its relationship to stochastic scenario testing.

MR. JOHN M. O'SULLIVAN: I'm the moderator and a speaker for one of the sections of this particular session. I work as a consulting actuary for ASA. My background is that I've been working with variable products since about 1982, and, like Frank Clapper, I'm a member of the Academy Committee on Variable Annuity Guaranteed Living Benefits (VAGLB). We have Dana Tatro who is a product actuary with Manulife Financial, which is a leading writer of variable annuities. Guaranteed living benefits is a vital part of Manulife's variable annuities strategy.

Frank Clapper is vice-president and corporate actuary of AXA Re Life Insurance. AXA Re has been quite active in the reinsurance of variable annuity guaranteed living benefits and guaranteed death benefits.

MR. CHARLES DANA TATRO: We have this exciting topic of living benefits and death benefits and the different guarantees on our separate account products. What I'd like to do is talk a little bit about the type of designs that we see on variable annuities. We might go into a little industry overview of the type of benefits that are offered on our leading companies out there. Then we're going to really get into the guaranteed minimum income benefit (GMIB). That seems to be a hot topic these days, and I don't think everyone really understands exactly what a GMIB is and how it works. Then we'll get into some state regulatory issues. These issues created a real hot topic due to California and New York not approving living benefits and the worry about some other states coming on board and jumping on the bandwagon.

Let's go through a few definitions of guaranteed benefits that are offered on variable annuities. There are basically four types. The first one has been around since the dawn of annuities and really came into age in 1994. That's the guaranteed minimum death benefit. These usually provide an enhanced value upon death, usually in the form of an annual step-up, some roll-up percentage on the client's investment or some type of ratcheting amount. These benefits have been around for a long time, and it is really no big deal in the industry anymore. Most of our companies that offer variable annuities have these benefits in the market.

The second type of guarantee that we see out there is fairly new to the market. These are guaranteed minimum accumulation benefits (GMABs), which are benefits that provide a minimum floor for your account value after a waiting period. These benefits aren't very prevalent in the U.S. If we do see ones in the U.S., the company usually controls the investor's fund selection. They're more prevalent in Canada, mainly because of the 75% guarantee regulation that they have on their annuity products. We don't think we're going to see too many more of these coming out, although it is something that a lot of companies are looking into.

The third type of benefit that we see out there would be the guaranteed minimum annuity floor, and that is a floor on a variable annuity payout. Essentially, when somebody does a variable annuitization, the company guarantees that their monthly or annual payout will never be less than some amount.

The final benefit that we have is the guaranteed minimum income benefit, and this benefit is a really hot topic in the industry today. This benefit provides a minimum floor of income upon annuitization. That is different from the guaranteed minimum annuity floor because this benefit is determined at issue, and it's more of a benefit on the amount that gets annuitized. The annuity floor, whatever amount you have at the time of annuitization, is what you get to annuitize. They just guarantee the payments thereafter. This benefit is exercised when the contract annuitizes. We're seeing a lot of them in the market today. It's in the form of a roll-up or a bull and bear, which is the greater of a roll-up percentage or an annual ratchet or some type of return of premium.

To kind of give you an idea of what we're seeing in the industry, I looked at the top ten companies and contracts that are being offered out there today, and out of those top ten, every one of them has a death benefit on it. Seven of them have a GMIB option of some sort on it, and one of them has the account value benefit that we talked about, or the GMAB. Obviously, the GMIB is really kind of the next phase of guarantees that we're seeing in the annuity industry, and that's going to be the focus of my conversation today.

At the time of annuitization, the client gets the benefit of annuitizing this fictional value called an income base at guaranteed purchase rates. This income base is usually based off the client's initial investment, and it might grow at some percentage or be ratcheted up by the highest anniversary value or just a plain principal guarantee on the amount that they have. This amount is applied to guaranteed purchase rates to determine what their floor of income will be when the client annuitizes after a waiting period.

As an example of how these things work, let's take a 6% roll-up GMIB with a seven-year wait for a male age 67 and a life ten-year period certain. The client puts in $100,000 and the account

value only grows to $130,000. The income base, which would be the 6% compounded over seven years on the $100,000, is about $150,000. The difference between the account value and the income base is about 15%. This is where the little catch comes in. The income base can only be annuitized at guaranteed rates. Those rates are lower than current annuity rates that are usually available at the time of annuitization. When you apply that income base to guaranteed rates, you get a value of 97 or $9,700. If you take current income rates, and these rates were done at 6% with standard mortality assumptions, you see that current annuity rates applied to the lower account value actually gives you a greater income amount on annuitization. Even though it looks like this contract is in the money, it's not. It's out of the money, and the client does not receive the benefit of annuitizing their GMIB feature.

Let's take the same example again. We assume that the account value only grows to $115,000. You have about a 30% difference, and you finally see that the GMIB income is slightly higher than the current annuity income. Like I said, that's the real catch. The difference between the guaranteed rates that they annuitize at under this GMIB feature and the current rate provides a lot of protection against any equity market downturns and equity market performance that we see.

There are several risk features to the GMIB that come into the pricing side that we look at when we're designing these benefits. The first, of course, is a roll-up percentage. That is what a lot of clients focus on, and it's the subject of a lot of compliance issues and disclosure issues. The higher that percentage is, the more risky the benefit. The second provision is whether or not you have a reset option on your contract. There are two, standard, reset options that we see in the industry today. One is an automatic reset that automatically steps up this income base to the contract anniversary value, and the other one is an optional reset. The difference between the two is in the optional reset, the client has to elect it. When they elect it, it extends the waiting period on their benefit from that time of election, where it automatically gives the client the benefit of locking in the highest anniversary value. The main difference between the two is after the waiting period is up, the automatic reset essentially gives the client a one-year put option on the equity market. For some reason, reinsurers love the automatic option better than the optional reset. If you're designing one, you might want to keep that in mind.

The next big provision that we see that affects GMIB pricing, which actually affects all guaranteed benefit pricing, is how you handle partial withdrawals from your contract. In the period of 1990–94, most of the withdrawal provisions allowed dollar-for-dollar withdrawal reduction on whatever benefit they were offering, whether it was a death benefit, an income benefit, or whatever. That is actually a very risky feature. Let's say you have a $100,000 benefit base and your account value is also $100,000. Two days later, your account value drops to $50,000 and your benefit base is still $100,000. With dollar-for-dollar withdrawals, your client can withdraw $49,000, leave $1,000 in his contract, and still have a benefit amount of $51,000 because you just reduce it for that withdrawal. That's pretty cheap insurance in anybody's book. Under the proportional withdrawal language that is coming out now, you actually look at the percentage of the account value that is being withdrawn and withdraw that same percentage from your benefit base. With the same example, that would have been a 98% reduction in account value, from $50,000 to $1,000. You would have had a 98% reduction in your benefit base, from $100,000 to $2,000. There is a big difference.

The next piece that comes into play here is your assumed interest rates on annuitization for your guaranteed purchase rates. Obviously, the better you make your guaranteed purchase rates, the less of a difference you're going to have between your guaranteed rates and your current rates. That means there's less of a market performance adjustment that you could absorb before your benefit gets in the money.

The next side is the deferral period. We see a lot of deferral periods between seven and ten years. The longer you go out with the deferral period, the more time you have to adjust to any market performances or one-time market shocks. It reduces the risk and cost of this benefit. The big unknown here is your annuitization election rates. You can offer this benefit to all your clients, but if none of them elect to annuitize, the benefit is not going to cost you anything, and that's really a big subject of debate when you do your reserving or your capital requirements or even your pricing. How many people do you assume actually annuitize when this benefit is in the money? As we all know, annuitization is still probably the longest four-letter word that we see out there in the annuity business today.

If we look at some of the products that are offered in the industry today, we usually see a deferral period of between seven and ten years. Some go out as long as 15. Some might even adjust it based on the annuitant's age at the time of purchase. Probably the most prevalent one that we see is ten years. We see income base growth rates of anywhere between 3% and 7%. We see benefit costs ranging from 25 basis points per annum to 50 basis points per annum, and those are based off of the actual income base, not the account value. There's one company out there that I know of that does base it off the account value. Most of them contain the automatic reset feature mainly because it's easier to hedge against. That's why reinsurers like it so much, and most of them will contain proportional withdrawal language in order to protect against that policyholder antiselection I spoke of earlier.

I have some general observations about filing these contracts. Thirty states usually won't give you a problem with them. There are a few that just rubber stamp them; unfortunately, I've yet to see too many companies that have a lot of sales in one of these states where they could actually just get it in there and rely on that state to generate some volume for them. Nineteen states will approve these with some significant questions. Two states, California and New York, will not approve these at all. We're going to talk about California and New York later. Most of the states don't have any guidelines to follow. You submit your product, and they'll come back with the questions.

Some of the more common state issues that I've heard about in my filing of GMIB is that it's inseparable from the base contract. Companies will add this benefit in the form of an optional rider. Some states come back, and they say that it's too major of a benefit. It should be included in the base contract. If any of you have ever filed a product, and you've had to refile the entire contract, you know that that could be quite an undertaking. Riders are much quicker to get through the states. You really want to fight against any state that says you have to include it in your base contract language.

New York is big on application disclosure. The other states are starting to follow their lead. Consumer disclosure requirements are needed because they indicate how you're going to sell this. They want some marketing material. One state out there has come back and said that's

discriminatory because you're not offering it on all of your annuity products. One state came back and said, you can't apply this to fixed account options. They say that what you're protecting against is variable fund performance. Why should you charge your client if they're in a fixed account option and decide to elect this benefit? That can be very difficult to do from a systems side, and not so much from a contract side. You have to be careful there as well.

Some additional issues that we see are that some states say that this benefit is subject to standard nonforfeiture law. They'd like to see cash-flow testing on it. You need to do quite a bit of work there. Some states like to see self-support tests and adequacy pricing, which kind of gets into the realm of your compensation to the general account for accepting this type of risk. These states don't have any guidelines as to how to do that self-support pricing. You're going to spend a lot of hours with their actuary going over what you did. They're concerned about reserves and whether or not you're going to hedge this benefit or whether or not you're using reinsurance. Oftentimes, if you tell them that you're going to reinsure it, they want to know what the reinsurer's going to do with it. That's a change that came about in the last couple of years. It used to be if you were reinsuring a death benefit, they said fine, no problem. If you talk about that with the income benefit, they get a little more scared. They want to know what the reinsurer is doing with the risk as well.

One of the problem states that we see is California. It pulled any filing of GMIBs back on May 30. Minnesota, New Jersey, New York, Texas, and Washington are other problem states. Those are really the big six that you have to worry about, and, if you look at the names in those states, they're usually the ones with the biggest sales. We're going to delve a little bit into New York because they're kind of their own separate entity anyway. They have yet to approve a GMIB. In discussions with the department, I have discovered that some of the actuaries feel that it's a two-tiered annuity, and they don't like it. They feel that's subject to standard nonforfeiture law. If you have any roll-up percentage that's greater than 2.99%, you'll have to hold standard nonforfeiture values on it. They're concerned about how they're going to reserve and how you're going to mitigate the risk of this benefit. The last I heard is they're not very happy with the Academy proposal, and I know they're not happy with the keel method. Even if that goes through, I don't think you'll see New York passing it any time soon.

On the California issues, they did pull all approvals of these benefits on May 30. The California regulators have made some significant strides in what they were saying before. California's original stance was that they didn't like reinsurance of this benefit. They said that you could not separate this benefit from the contract and only reinsure that portion of it. It would not qualify for automatic reserve credit. They had stringent rules about reserving for your actuarial memorandum and your annual statement. They wanted a lot of guidelines on that. Since I wrote my slides for this presentation a few months ago, legislation has been passed in California. It is now sitting on the governor's desk. That legislation outlines what needs to be done in order to have a guaranteed living benefit available in California. The legislation requires the department to issue a bulletin explaining exactly what is needed for companies who had a benefit available in California and then had it pulled by the department. That bulletin has been distributed to a bunch of companies for review. By the end of September 2000, companies that can certify compliance with this bulletin will be able to issue their IB in California once again. That was a major change. That's it for my proposal. I withdrew this for you guys.

MR. O'SULLIVAN: Dana gave some background from a product development actuary's viewpoint, and the state resistance on reinsurance and reserves. I'm going to spend my time focusing in on the proposed actuarial guideline. By way of background, to kind of put the regulatory setting into sort of a framework, we started with AG 33 which said that, in setting the reserves for a variable annuity, or any annuity, you would consider all the possible benefit streams. You have to take into account all the guarantees. AG 34, which was MMM, takes care of the reserves for the variable guaranteed minimum death benefit. The proposed guideline quad M, or MMMM, deals with guaranteed living benefits.

There are some similarities, and there's some differences with AG 34. Similar to AG 34, the proposal goes ahead and uses an integrated Commissioner's Annuity Reserve Valuation Method (CARVM) calculation, and we'll talk about that a little bit more in a few minutes. You calculate a reserve that incorporates all the guarantees in your contract. Like AG 34, there were the same five asset classes, the same five asset classes. Unlike AG 34, though, there's no drop in return. Instead, what is basically used is a simplified stochastic method.

The basic structure is you start calculating the integrated benefit reserve, and, as I mentioned, that includes all the guarantees in the contract. You have to do it including the guaranteed minimum death benefit and including any guaranteed living benefits. Then one calculates what the reserve would be without any of these guarantees. That has been termed separate account reserve. It has an intuitive feel, but it might not all be in a separate account. You could have a fixed account or a market value adjusted account (MVA) or something like that. The separate account reserve is sort of a useful tag, but it doesn't have its usual connotation. The excess of the first over the second is the extra reserve you're going to hold because of the guarantees. It's basically a "with" and "without" parallel calculation. You do it "with," and you do it "without."

Let's focus in on the "with" part of the calculation and the streams that involve guaranteed living benefits. First, you must project future investment performance. That gives you a projected account value. You know what your guaranteed living benefit would be at that duration. The difference between the projected account value and the guaranteed living benefit really becomes a net amount at risk. Working with that net amount at risk basically gives you that part of the calculation that deals with the payment of the guaranteed living benefits. Nothing about this subject is terribly simple. Getting policy forms and approvals is not simple. Reinsurance is not simple. Reserve calculations are not simple. When you come to apply this net amount at risk concept, it's not simple either because, depending on the kind of benefits that you have guaranteed, the net amount at risk is defined a little bit differently.

I'll give you an example. If you're dealing with a guaranteed minimum income benefit, sort of like what Dana had shown, you come up with the present value of the income payments that you're guaranteeing, and you compare that against your account value. That difference becomes your net amount at risk. If you're dealing with something like a guaranteed partial withdrawal feature or a guaranteed minimum payment floor, in a variable immediate annuity, then what you're really looking at is what the income payments would be with and without the guarantee to get what your net amount at risk is.

I guess this is a good time to kind of mention two things to you. One of them is that a lot of my background has to do with variable immediate annuities, and the area that I've been

concentrating on in the Academy group has been this whole area of payment floors, among other things. This particular guideline would apply not only to deferred annuities but also to immediate annuities, and to the payment floors on those things.

The second thing I need to go ahead and tell you about is that you really need to be careful in applying this kind of guideline. There are a bunch of principles that are set out, and you really need to understand the principles. You always need to use actuarial judgment that the stuff is applicable to this particular benefit or think carefully as to whether or not it is applicable to a benefit. For example, the wording of guaranteed living benefits is quite broad, as I mentioned before. For example, if you were to have some sort of a persistency device in your variable deferred annuity that guaranteed an annuitization bonus if you didn't hit a certain bogey in ten years, or an investment performance, or provided for a reduction in future annuity charges, that would be swept into a guaranteed living benefit. You'd need to look at the principles, and at the guideline to see how you would apply it to that particular circumstance.

We have the integrated benefit reserve. The next issue is, how do we go ahead and calculate these interest rates that we're going to use to project the account value? There are three methods. One of them is sort of borderline insane. It uses a full stochastic method on your valuation. If you had three or four months to do the calculations with a good block of business, that might be an option. I'm being a little bit sarcastic. There are probably smarter people out there who could do it in a reasonable amount of time. But that leads you to two other choices. One of them is called the keel method, and we'll get into talking about that in a minute. It is what's termed a safe harbor. The second one is to come up with basically your own keel method. In other words, you do the work paralleling what the Academy had done and come up with representative scenarios. It might be 10–12 scenarios. You do all of the same kind of analytical work to show that this is a reasonable approximation to go ahead and come up with a reserve for this particular guaranteed living benefit.

The safe harbor or the keel method is not mandated. Depending on the design you have, using the keel might produce excessive reserves for you. You shouldn't use it blindly. Second, you might not qualify to use the keel method. If you have a benefit that's path dependent, like an

annual reset or something of that nature, then you can't determine what that future guaranteed living benefit is, and you can't use the keel method. So, you have to then use either full stochastic or the representative scenario method.

The keel method is a standard formula. The term keel comes from the fact that it looks like the keel of a boat. It drops down and then kind of comes back up. If you were to graph these interest rates, the graph looks like a keel. Everybody kind of struggles with the keel, and everybody has their own way of looking at this thing to make sense of it. It is a statistical device to go ahead and come up with cumulative return for various periods of time. Once you have those returns, then you can go ahead and use them to project the net amount at risk, and the same formula gets used for each of the asset classes, but you would use different means and standard deviations.

These are the means and standard deviations for the five asset classes. You'll see this in the September report of the Academy group on VAGLB. The June report is out on the Academy website, and I believe the September report is the one that got presented to the Life and Health Actuarial Task Force (LHATF) recently. I urge you to study them. If you have any questions talk to any of the people that are on the committee.

Table 1 shows the keel parameters. There are the means and the standard deviations that one would use. The actual keel mechanics use a formula that is based on a lognormal distribution. It's a classic formula you can find in a textbook like Hull. We're using the 83.33 percentile because you have to figure out what the probability is that you want to put in there. Setting the probability at the 83.33 percentile, is something that LHATF has to confirm. That's the appropriate percentile to be using. These are gross rates. You have to adjust for your actual fund charges, and then the actual product charges that you have and both the basic annuity charge and any VAGLB charge that you have.

TABLE 1
Keel Parameters

Asset Class	Mean	Standard Deviation
Equity	13.18%	12.70%
Bond	9.14	7.05
Balanced	11.03	9.48
Money Market	7.54	2.70
Specialty	11.95	13.03

On the representative scenarios, you would develop your own set of rates to project future account values. There's a requirement to calibrate it against the Academy work. You would need to use a lognormal distribution. You have to use the same mean and standard deviation, and it has to be sufficient at the 83.33 percentile. You have to find the scenarios that are both a good approximation and that are workable for you. If you don't use the keel, then you have to develop the alternative scenarios. You have an annual certification that they are appropriate. You have to maintain documentation in your files so insurance regulators can review, if need be. Then you have the ongoing monitoring that it's really still appropriate. You're going to simplify things for these representative scenarios. Hopefully, you've caught all the major elements, but it's conceivable that one of the major elements that's going to drive the risk on this thing wasn't incorporated into the representative scenarios. If that is the case, then you have to go back and redo your work.

I'm going to show a simple example. I did a guaranteed minimum accumulation benefit. As Dana said, it's not the most popular benefit, but it's the easiest one to use to show you how things work because it reflects one point in time. My simple deferred annuity has no surrender charges. It has a death benefit equal to the account value. It has an accumulation benefit with a 10-year waiting period. If your account value is below your premium, we'll top you off at the end of the 10th contract year. The charges that I assumed in the example were 135 basis points for the basic annuity and 40 basis points for the GMAB. My fund expense assumption was 70 basis points. With these assumptions, you should be able to go through and duplicate the results of this particular example.

I'm going to look at a case of a 65-year-old male who has a $50,000 single premium. There's no withdrawal activity on the contract. All the dollars are invested in an equity class. My valuation date is at the end of year seven, in contract year seven. So, I have only three years to go to the exercise date. The account value is $40,000. For the equity class, after I adjust for the fund charges and the product charges, my cumulative returns are: -1.54% (year 1), 4.17% (year 2) and 11.53% (year 3). I'm assuming that my valuation year and my issue year match up. I'm only dealing with a one-year return, a two-year return, and a three-year return, and these are cumulative returns.

The first thing to do is figure out the integrated benefit reserve. I'm going to look at the benefit stream that involves a full surrender at the end of year ten, which will give you the greatest present value in this particular set of circumstances. I'm going to calculate the value of three benefits. One of them is the shortfall that has to be provided by the GMAB. Another one is the basic surrender benefit to the survivors, and then the death benefits that will be paid in years eight through ten.

This GMAB benefit is a return of premium. The premium was 50,000. That's where I got that from. If you apply the three-year cumulative keel to the $40,000, you would get the projected account value. The difference between those two is the net amount at risk. Then, if you present value it for interest and survivorship, you come up with $4,186. One can combine the present value of the death benefits, the present value of the basic surrender value, and the present value of the GMAB shortfall. The integrated benefit reserve would provide for all the guaranteed benefits in the contract, and that's $42,306. The difference between that and the $40,000, which was the account value, is what we're holding in a general account for the guarantees that we've made. That turns out to be 6% of account value. That's a real simple kind of example you can take home. If you have an interest, you can duplicate all these numbers out of either the June report or the September report, in the appendices.

Let's put this stuff into perspective. One of the more subtle points about the way that this guideline works is you have a pattern where, when you get closer to the benefit exercise date or you get more in the money, you have to set up more and more reserves. As such, you should not

lull yourself into looking at just what happens when you're at the money because it doesn't take much going into the money to drive up these reserves quite dramatically. One of the points that Dana showed was for the guaranteed minimum income benefit. When it came down to 30%, you were suddenly right in the money.

In Table 2, we calculated a return of premium GMAB for valuation date seven, assuming that our account value was 40,000, which was 80% of what the GMAB benefit was. That reserve came up to 6%. You can see that at duration five, or at the end of contract year five, we would not have set up any kind of reserve. As we get more in the money or the duration gets closer to the exercise date, those reserves can really balloon on you.

TABLE 2
Sample for Return Premium

Valuation Date	Asset Value as Percentage of GMAB		
	100%	90%	80%
5	0%	0%	0%
7	0	0	6
9	0	10	23

If we look at a 3% roll-up, as in Table 3, we see a similar kind of pattern. Because the benefits are more, the reserves are more. You should remember that these are the minimum statutory reserves that you have to hold. You could hold more than these, and there are devices that can be used to smooth out the impact on earnings of this kind of volatility or fluctuations in the reserves.

TABLE 3
Sample for 3% Roll-Up

Valuation Date	Asset Value as Percentage of GMAB (on Valuation Date)		
	100%	90%	80%
5	0%	0%	2%
7	0	3	15
9	3	13	26

What is the status? The guideline has been proposed. We need to do some more revisions on this thing. The most likely guess is it will apply to the end of the year 2001. It will apply to all policies issued since 1981. As with AG 34, there would be a three-year grade-in with regulatory approval. One of the things that I'd like to solicit your input on is that as we were drafting the guideline, we had a lot of examples and sort of other good stuff in there. What we decided was that the guideline was very long to start off with, so that we would put a lot of this stuff into a Practice Note where it would be more appropriate.

For example, we're going to give a lot more guidance on representative scenarios. How do you do the sort of proving to yourself that it's a good approximation to full stochastic and things like that? If you have any ideas or stuff that you want to see, please let us know. What we're hoping to do in this next quarter is to really get most of the Practice Note done. If you can help us or if you have any questions, please don't be bashful. We can use all the input we can get.

Frank has some really good stuff on reinsurance, risk-based capital (RBC), what's happening in Canada, and on different reserve techniques that we could go ahead and use to smooth out the earnings.

MR. FRANKLIN C. CLAPPER, JR.: First, I want to illustrate how the product behaves graphically, showing the benefit patterns and the premium patterns and so on. This will bring out the kind of risks we're dealing with from a valuation and reinsurance point of view. Then I'll talk about how reinsurance is done and how it's valued. I hope to be able to talk about what I'm calling a generalized valuation method, which is really a loss-ratio approach. I think it's a little more satisfactory for GAAP than the deterministic method that's used in the guideline.

Variable annuity guaranteed benefits are very different from what valuation actuaries are used to dealing with. They have a low expected claims rate but a very high standard deviation. That's quite different from most life products, and that's true for both death benefits and living benefits. The main risk is financial rather than insurance. Therefore, we have to borrow from other areas of practice. We're not really in a regular, traditional life valuation area. We have to think about financial techniques and nonlife methods.

Chart 1 is a guaranteed minimum death benefit (GMDB) with a simple return of premium. The level line reflects the guaranteed benefit by policy year, and the jagged line shows the variable benefit by policy year, and, of course, when the guarantee is higher than the variable, you have an amount at risk. You can illustrate this for other benefits.

Chart 2 shows the ratchet. The ratchet can only go up. It can't go down. You have the guaranteed benefit going up to here and staying level and then climbing up. Meanwhile, the account value is dipping down. Then it comes back up. The ratchet has a different benefit pattern.

Chart 3 is a roll-up. You have the smooth curve for the guaranteed benefit. Where the variable benefit dips below the guarantee, you have exposure.

Chart 4 is what we call a combination. I guess this is what Dana calls a bull and bear benefit. That's a little bit richer than either the roll-up or the ratchet. Those are the benefit patterns. Chart 5 shows the premiums and the benefits each year under one scenario. I just picked a random scenario. It's probably a worse-than-average scenario. The square line reflects the benefits that start out at zero and go through some movement. This is the death benefit. The revenue (diamond line) meanwhile is going more or less flat. We have some exposure.

Chart 6 shows the difference between the revenue and the benefits, or the profits, without any reserves. First, the lower line shows the year-by-year result, and then the upper line shows the cumulative result. The point is that we'd like to smooth these things out somehow both from a risk point of view and from an accounting point of view.

The GMAB, the accumulation benefit, of course, has a very different pattern because all the benefits are paid after the window, if at all (Charts 7 and 8). There are no benefits until year 10. The same is true of year-by-year profits and cumulative profits. The profit ends up being positive but not by much.

This just helps understand the risk structure that we're dealing with for reinsurance purposes and for accounting. I'd like to tell you a little bit about reinsurance, since I'm the only reinsurer on the panel. I also want to talk about risk transfer because it is an issue in some jurisdictions. It turns out that the way CARVM works is not obvious on this product, at least not from the reinsurer's assumed point of view. In many cases, the reinsurance is not done on a strictly proportional basis. You have to consider how that affects valuation as well.

When you have separate accounts involved, you really can't do coinsurance. Even if you had a separate account, it wouldn't be the same as the ceding company's separate account. You can do modified coinsurance by setting the modco interest rate equal to the earned rate on the separate account. The reinsurer would pay an allowance, and so forth, and the reinsurer would fully participate in the entire policy.

But the usual form of reinsurance is what I call risk premium reinsurance. It's a little bit like yearly renewable term (YRT) but not exactly. Only the guaranteed benefit is reinsured. The reinsurer doesn't participate in the separate account. The reinsured benefit is defined as the net amount at risk. In other words, there's only a reinsurance benefit if the guaranteed benefit exceeds the nonguaranteed benefit. The reinsurance premiums are based on some combination of the account value and the guaranteed value. They might be age-adjusted and sex-adjusted for mortality. They might take into account the net amount at risk, as I said. Sometimes they're adjusted for year-by-year claim experience. They can go up or down. There's a minimum premium and a maximum premium. The minimum premium involves both a minimum rate and a minimum aggregate amount. That's to protect the reinsurer when it otherwise wouldn't collect a premium. Likewise, there's a maximum rate that protects the ceding company so that it is not totally exposed to premium increases.

To illustrate how reinsurance actually works graphically, Chart 9 shows a GMIB benefit. In this case, of course, you don't have any benefits until you get past the window for annuitization. The benefits depend on the annuitization experience. In this scenario, they climb up and they go back down. Meanwhile, we have fairly smooth revenue.

Chart 10 shows results for the ceding company, net of a 100% reinsurance deal. There are no benefits because the reinsurer pays them all. The reinsurer takes a piece of the revenue, but it's proportional so that the ceding company is still left with a relatively smooth revenue stream.

Chart 11 shows the year-by-year profits and the cumulative profits; then Chart 12 shows the same thing after reinsurance. The year-by-year profits are fairly level, and the cumulative profits just keep going up.

Reinsurance, as I said before, might have some nonproportional features. It doesn't have to, but it often does. In particular, if the ceding company wants to make the reinsurance cheaper, there can be a deductible that simply means that the ceding company pays first-dollar claims up to a certain aggregate amount. Usually it's expressed as basis points times the in-force account value. Likewise, the ceding company might elect to keep part of the tail risk and self-insure on the stop-loss aspect of this, in which case there'd be an aggregate claim limit, above which the reinsurance would not pay. There can be per-year annuitization limits; for example, the reinsurance would only pay up to 30–35% or something like that per year. There are also forms of reinsurance that are pure stop-loss.

In regards to reinsurance, risk transfer is an issue for some regulators, especially for risk premium reinsurance. The Reinsurance Model Regulation says that all significant risks must be transferred in order to qualify for "reinsurance accounting" treatment. The purpose of the regulation was to avoid proportional reserve credit when a reinsurance arrangement does not transfer proportional risk. It seems clear that you wouldn't be entitled to proportional credit if the arrangement is not completely proportional. Normal coinsurance, modified coinsurance, and comodco, if they transfer all the significant risks, are proportional, and they qualify for proportional reserve credit. However, note that the regulation does not disallow nonproportional reinsurance. Many people misinterpret that. In fact, the regulation specifically excludes YRT, as well as all forms of nonproportional reinsurance, such as stop-loss. It doesn't really comment on how those are supposed to work.

Traditional YRT, which covers only the mortality of risk on traditional policies, is a nonproportional form of reinsurance, but the ceding company takes proportional reserve credit with respect to the mortality risk only. Only the mortality reserve is considered. Coverage and reserve credit are proportional, but not on the whole policy.

Risk premium reinsurance of variable annuity guarantees, whether it's on death benefits or living benefits, is very similar to YRT as long as premium amounts are capped so that the ceding company is not exposed to the complete financial risk. The financial risk has to be transferred to the reinsurer.

I should point out that there aren't any formulas for valuing pure nonproportional coverages because they don't happen that often in a life business. They're mostly in the nonlife business. I talked to some of my nonlife colleagues, and they don't have any strict rules pertaining to how to do stop-loss. Normally the ceding company would get no reserve credit, though, for something like that. The bottom line is that if you have a typical risk premium reinsurance arrangement, it often has a nonproportional element to it. You have to consider that. Both the reinsurer and the ceding company need to modify the reserve credit, and that's a judgmental aspect of the valuation, both for statutory and for GAAP.

How is reinsurance valued? It is valued in a straightforward manner for coinsurance or modco. The risk premium reinsurance on death benefits is covered by Guideline 34. Risk premium on living benefits is incorporated into the guideline that John talked about. There are not a lot of rules. It mostly says the same thing as Guideline 34, and it does comment that nonproportional aspects need to be considered.

In regards to risk transfer in CARVM, I want to point out that the ceding company and the reinsurer are in different positions. That's what I meant when I said it's not proportional. Even under full coinsurance, the reinsurer might not invest exactly the same way as the ceding company. The reinsurer might pay an allowance that covers the ceding company's expenses, but they're not exactly the same. There really isn't any such thing as pure proportional reinsurance. When you get into deferred annuities, there's also a difference in the optionality. The reinsurer is

participating in only part of the contract, not the whole contract. The reinsurer typically is not even participating in ceding company risk, such as normal annuitization or cash surrender, for example. They're really in different positions.

This affects how you value the contract, especially on the assumed side. The committee actually discussed four different theories. One is "mirror reserving:" the ceding company takes credit for the same reserve that the reinsurer sets up. The second theory is the so-called "dependent-path" method: the reinsurer holds a reserve on its portion of the risk using the same optionality assumptions as the ceding company. Under the third theory, the "independent path," or the reinsurer, ignores what the ceding company is doing, and it holds a full CARVM reserve, which is the greatest present value for the risks that it has assumed. We didn't think that was necessarily appropriate. The last theory, *estimated utilization*, means using option probabilities, and, of course, regulators aren't too fond of that on annuities. I'd like to talk a little bit about a reserve method that I thought you might find helpful. It's different from what is in the statutory guideline. There is a problem that we're faced with in this class of products. There are several issues. One is that the revenue and the benefits are not only not correlated, but they're somewhat negatively correlated. Revenue goes up just when benefits are lower or vice versa. When your experience is bad, the benefits go up and the revenue goes down. That's a little different from some of the other products. The premium, as I said before, consists mainly of a margin for deviation as opposed to an expected result. That's more of a nonlife concept than it is a life concept.

Even if your basic pricing assumptions are met, the benefit pattern can be very volatile from period to period. You don't really know what's happening until it's too late. Likewise, there's a basic assumption in some of our reserving techniques that we have some mean reversion. In other words, if we have a bad year, we expect to go to a good year to sort of balance things out, but with this product that's less clear than usual.

When you're valuing variable annuity guarantees, there are several things to think about in setting reserves. The main one is what should your net-to-gross ratio be? It's fairly easy to project your revenues, but you really have no idea what your benefits will be relative to those revenues. Some people have used a retrospective calculation where you assume that your pricing

is correct, and you accumulate premiums minus benefits and hope that it's enough. Obviously, you don't know whether it's enough or not. Therefore, you have to do a prospective calculation as well. You really don't know how to interpret experience because there's so much random fluctuation from year to year. And all this comes together to create a hideous GAAP earnings pattern if you're not careful. This is the problem I was trying to solve.

If you do hold a conservative reserve in the beginning, you don't know how fast you need to release it. In other words, should we keep it for future experience or not? There are different ways to approach this. Obviously, the most complete and accurate way would be to do complete stochastic testing every time you want to do a valuation, but that's a little bit onerous. It does recognize all the product features and contingencies, but it's still not absolutely accurate because your assumptions are still subject to question. It takes a lot of time to do this and requires a lot of judgment, even though you're doing some accurate calculations.

The retrospective method is easy to understand. You can control your earnings patterns, at least in the early days, but you don't know whether it's enough, and you don't know when to let go of it.

The statutory methods in the guidelines are what we call deterministic prospective methods. They're designed to imitate some percentile of scenarios, as John mentioned. They're easy to understand. They're pretty easy to calculate once you get used to doing it. They do adjust the reserve for moving in and out of the money. They will be sufficient if your experience is close to what you expect but they can be very volatile. If you calculate these things from quarter to quarter or year to year, they go up and down according to whether you're in the money or out of the money. If things go bad, they can be insufficient. You have to start building up the reserve, and by then, it might be a little bit too late. They might not be applicable to every situation. In other words, the Guideline assumptions can't always be used for every product. Some products don't fit into this framework.

I came up with something based on a rearrangement of old concepts. I'm calling it the retrospective/prospective method. It's really a self-adjusting loss-ratio method. You can

think of it in terms of *FAS 97* DAC amortization. That's really where the idea came from. You project benefits and premiums "as desired." That means the method doesn't incorporate any assumption about how to do that. That's not part of the reserve method. That's part of your conservatism and your outlook. You can do it stochastically and pick a scenario. You can do it deterministically with very conservative assumptions or however you think it is best. That's not an intrinsic part of the method. You have a prospective stream of benefits and premiums and you have a retrospective history of benefits and premiums. You force the retrospective reserve and the prospective reserve to be equal by adjusting the net-to-gross ratio. As I said before, it's like *FAS 97*. It's self-adjusting.

In practice, you project benefits and revenue each year. You don't have to redo the projection of this for every valuation date, but you have to think about it. At each valuation date, you replace your expected benefits and premiums for the period that just passed with the actual, and you recalculate the net-to-gross ratio from inception of the block that you're valuing. This calculation is done at an aggregate level. We do it for each reinsurance treaty separately.

You calculate the K ratio. That's the part that's not in your handout. It's basically the ratio of past and future benefits to past and future gross premiums. I don't allow it to exceed 100% because I don't think the reserves should be less than the gross premium reserve at any time. The final reserve is the present value of benefits less your present value of net premiums, being the K ratio times the present value of gross premiums according to your projection.

As I said, the projection can use a deterministic or a stochastic method. You can change assumptions when you think it's appropriate, but you don't have to change your future assumptions every period. You'd use the same model. You might update the projection for in-force business, though. That's advisable because the in-force business changes rapidly. You're also adjusting for your in-the-money or out-of-the-money experience as your in-force business evolves. You can control the degree of conservatism. For GAAP, we're trying to maintain the same relative level of conservatism from period to period because we don't want to be accused of manipulating earnings.

I'd like to illustrate graphically how this works (Chart 13). I'm projecting the experience, including the reserve increase from the inception of the block. There are three elements in here. The benefits are reflected by the dark part of the bar. This is showing the relationship to premiums as a percentage. The reserve increase is the light part. The earnings are what's on top. The reserve increase doesn't build up much in the early durations because the benefits exceed the average ratio, but reserves do start building up later in order to anticipate the future benefit experience, where the reserves are released. We get an earnings pattern that is as smooth as possible, and this is on a projected basis.

Chart 14 shows an actual history of reserves and earnings under one scenario. There's more variation here because actual experience is more volatile than projected experience.

The benefits start out at zero, jump up, go back down, and then they really go up. The reserve kind of fills in the blanks. It increases, and then it decreases a little bit; it increases some more, and finally it's released at the end. We don't have exactly a level percentage of earnings, but at least we have earnings in every year, which is a major accomplishment on this product.

The advantage of this method is we try to make the reserve change as smoothly as possible. We adjust our experience as soon as we know what it is, and we spread that experience over the past and the future. We also recalibrate to the actual in-force business to account for changes in the distribution of the in-force business by age, sex, and so forth, and degree of in-the-money. In other words, the account values are in or out of the money according to the actual financial experience. If you like, you can revise the projection to adjust assumptions at the same time.

You true up to the actual experience as quickly as possible, but you keep your conservatism with respect to future experience so that it's consistent. Reserves are released as quickly as possible without overdoing it. You still have some set aside for the future. Another way of looking at it is that your sufficiency criteria is trued up as you go ahead.

I'd like to cover just a few odds and ends. Many people talk about how the risk is managed on this product. At AXA Re we take a very broad view of risk management. It really starts with

product design, and we're careful about that. We actually help clients design products so that they're not too risky. Obviously, they need to be priced adequately using sophisticated modeling and underwriting. What do I mean by that? You don't usually underwrite annuities, but you do have to take a general underwriting approach to make sure you're not being exposed to some unexpected risk based on the product design or some combination of benefits where you exacerbate the risk or something like that.

The next phase of risk management really is oriented towards the data. You need sophisticated administration systems to keep track of it, to value it, and to model it. You have to do this on an ongoing basis to measure your risk. Finally, what people usually refer to as risk management is either reinsurance or hedging to get rid of the risk that you don't want to keep. That's basically it.

I was going to talk a little bit more about the advantages and disadvantages of reinsurance and hedging. Reinsurance obviously covers all the risks, including the financial risk. It's easier than hedging to set up and administer. The accounting and valuation is not totally clear, but it's relatively clear on many products. There have been comments in the press about reinsurance expense. Obviously, a reinsurer needs to be paid for its services. The hard cost, if you will, is going to be higher than hedging because, what is reinsurance except hedging plus a service fee. That's basically what it amounts to.

If you look at hedging by contrast, you're only covering the financial risk. You're not covering the insurance risk of the product, the mortality risk, or the annuitization risk, or those things that have to be estimated so that you can match your options on the asset and liability side. If you want to do hedging, it does require a significant operational investment that presumably would need to be spread over a large base in order to be effective. The accounting for hedging is not clear, even if you do it right, and the valuation is also not clear. There's no allowance for hedging on the statutory side, so you don't get any credit for it.

For risk-based capital, the living benefits have a somewhat onerous formula. It was done on an interim basis because they didn't have enough time to develop a really sophisticated formula, and

it's simply a factor, as usual, times the entire reserve, both the separate account and the general account. That factor is 3% if you don't do a Section 8 actuarial opinion. If you do an opinion, then it's still 2%, if a policy is in the money or 1% if it's out of the money.

This formula is supposed to be revised as a follow-up activity of the C-3 subgroup, and they're just starting to work on that now. They were waiting to see what the results were from the Academy committee with respect to reserves, but it will be worked on. It's the same basic group that worked on the first C-3 modification.

It is interesting to know, for professional and academic reasons, that Canada takes a completely different approach, as they usually do. They initially had a very onerous guideline, and that was simply trying to force the industry to come up with something better. The CIA did come up with a fairly comprehensive, albeit not final, guideline. It involves stochastic reserves but there is a formula for total reserves and capital because they didn't quite develop the full guidance for how to do that stochastically. This report can be found on the CIA website. It's not hard to find at all.

MR. KERRY A. KRANTZ: You mentioned estimated utilization isn't popular with regulators. I'm thinking about UVS. Could you comment briefly on what you think would happen in a future regulatory environment, assuming UVS is adopted? What is the point about estimated utilization?

MR. CLAPPER: The comment about estimated utilization was made on behalf of the committee. That whole slide was oriented towards the committee discussions. It's not a personal opinion. We considered all four methods, but we felt that it was our mandate to stick to CARVM. In that context, you can't use utilization. It was that simple. That's why I made that comment. Now, if you're going to use UVS, I have probably the same or less understanding of everybody else in the room, but I'll try to comment. I like the Canadian approach because it gives enough guidance so that you don't make mistakes. I mean that's basically what they do, and they keep building it up. If actuaries don't give a satisfactory result, then they try to impose further guidance until it becomes almost a rule book. There are degrees of doing this, instead of

doing it all at once. So, in that respect, there would probably be some guidance on appropriate utilization assumptions in that context.

MR. STEPHEN A.J. SEDLAK: I guess my question is what you seem to be coming down to is a bit inconsistent with AG 34, the guaranteed minimum death benefits. Is there any thought to retrofitting that? New York had a draft regulation, which was 185. The state still hasn't adopted it. It contained a whole bunch of things, one of which was AG 34 in regulatory form.

MR. CLAPPER: I can start on the first part. I'm not sure what you mean by inconsistent. It's viewed as being totally consistent. The reason we have the keel method is simply because you have a different risk pattern, and some of that has been discussed a lot. This brings out the fact that models for different products seem to be inconsistent. They do have to be biased towards the risk in the particular product that you're valuing. In this case, the risk pattern is completely different for living benefits than it is for death benefits; therefore, you get a different set of deterministic approximations to the stochastic model. The stochastic model underlying all this is completely the same between the GMDB and living benefits and all the products. It's the product itself that dictates the deterministic approximation to that background.

MR. O'SULLIVAN: The keel method is set up to really display a long period of under-performance, which is the reason that it's not a simple drop and rise type of thing. What one is doing is putting in the 83.33 percentile for each of the time periods involved to come up with a cumulative return. In some ways, the keel is a drop and a rise, but it's chosen so that it basically gives you a reflection of continued underperformance. With respect to New York's reaction to the Academy proposal, they had done some independent modeling in the late spring/early summer and were concerned with the pattern of reserves. In particular, they expected larger reserves in the early years. Under the current proposal, this VAGLB reserve is set up so that it's self-correcting over time, which can give you more fluctuation on the earnings pattern. I suspect the same idea of the fluctuation on the earnings is underlying some of California's concern about whether transfers of people moving funds around should be treated as an elective benefit. In other words, does this reserve standard reflect the fact that people can transfer their funds

adequately? If somebody was in a more conservative fund, and then moved over to a more aggressive fund, will the reserve be adequate?

MR. TATRO: I guess I'll comment on the New York question. I haven't heard anything.

MS. CINDY D. BARNARD: A reference was made to the AICPA coming out with a Standard of Practice on nontraditional products. I was wondering if it was going to reference GAAP issues for these types of VAGLBs and anything in relation to what Frank talked about.

MR. O'SULLIVAN: I know a little bit about it, but not a lot. Tom Campbell would be your best source for this thing. My understanding is that one would basically have to be setting aside some of the revenue that was charged for these guaranteed minimum death benefits or guaranteed living benefits. You'd have to be basically recognizing the possible benefit cost so you didn't prematurely release the charges into earnings. Rather, build up a reserve for possible benefit costs and release it over time, which has not always been the case on GAAP.

MR. CLAPPER: I did discuss the method here that I described to our auditors. My preliminary opinion is that it satisfies all GAAP principles. Depending on what the Derivative Investments Group (DIG) Committee comes out with regarding *FAS 133*, you could also handle hedging. What you do is incorporate the hedging cost in your projection and in the hedging benefits. What happens when you incorporate hedging in the reserve is that the reserve goes up because hedging is designed to cover the tail risk, not the expected cost. The benefits in an average reserve scenario will be less than the cost. That's one aspect of it. You can also incorporate DAC in this method. Just spread the cost proportionally to the revenue and that sort of thing. This method that I'm talking about has the advantage that the reserve change almost exactly offsets the difference between the actual and expected benefits in the early durations. The only other aspect is the adjustment in the net-to-gross ratio. If you think of it that way, it's a lot clearer.

FROM THE FLOOR: Your auditors gave you an opinion that this was satisfactory to all.

MR. CLAPPER: Not in writing.

FROM THE FLOOR: The second question is, is that true for a direct writer of these contracts or just from a reinsurer's perspective?

MR. CLAPPER: I'm not sure I can generalize it. Obviously, we're in a different position because we're looking only at the fees and the benefits on the net amount at risk. We're not looking at the whole policy. It might be a different story if we looked at the whole policy. I don't really know. I think there might be a difference. You might be able to follow the same principle if you calculated it. The person that I got the opinion from was the auditor's actuary, and we were just discussing it from a general GAAP point of view. We didn't really get into emerging guidance from the DIG Committee. If they come out with something different, then we'll have to change.