1998 VALUATION ACTUARY SYMPOSIUM PROCEEDINGS

SESSION 27TS

DEFERRED ANNUITY MODELING ISSUES

Meredith A. Ratajczak, Moderator

Francis P. Sabatini

DEFERRED ANNUITY MODELING ISSUES

MR. FRANCIS P. SABATINI: I'm Frank Sabatini with Ernst & Young. My co-speaker is Meredith Ratajczak from Milliman & Robertson. We're talking about deferred annuity modeling issues. I'm going to talk about variable annuity modeling, and Meredith is going to talk about fixed annuity modeling issues. Variable annuity modeling is our new challenge. The argument can be made, "Why model variable annuities when there is no risk involved? How many of you in the variable annuity business saw about a 10–15% reduction in income over the past couple of months? There might be a modeling issue in this, so my focus is on, the need for constructing good, sound models that allow you to do a variety of things, such as test, in particular, the sensitivity of earnings to moves in the markets. I am also going to talk about how important it is to do good, sound modeling as it relates to all of the new bells and whistles we're adding to products.

To set the stage, we want to talk about the product universe. The straight variable deferred annuity offers no risk and no pain. Equity markets grow, fee income grows, and production grows. The guaranteed minimum death benefit also is a well-established product. The remainder are fairly new. The guaranteed claim income benefit is basically the option to annuitize at a guaranteed benefit amount. Guaranteed maturity value is basically an underlying guarantee or return of premium on a variable product. Variable immediate payout annuities are self-descriptive, and pass-through products are variable annuity look-alikes in a general account contract.

We tend to think about modeling in four contexts: cash-flow testing, pricing, financial management, and risk management. From a framework perspective, there is no meaningful exposure to economic loss. You can't lose money on variable annuities, but your earnings can bounce around, and that's the major concern. With the other products, risk is in earnings volatility issues, unlike the traditional insurance product, primarily because they're linked to fund performance, which can be fairly volatile.

The first goal with respect to the goals of the modeling process, as it relates to variable-based products, is to represent fairly the level of sufficiency. The more rigorous the model is, the more likely you are to produce meaningful results that are consistent with the level of accuracy you want to produce in fixed product cash-flow testing. When interest rates go up for fixed products, it triggers certain options on both sides of the balance sheet that produce a different economic outcome than you would get if interest rates stayed level or went down. It's the same dynamic; the question is, how do you define it? Equity markets go up, income goes up, and lapses might go down. The same thing happens if the equity markets go down.

One of the difficulties many companies struggle with when testing fixed business under the standard seven scenarios, is determining which corresponding equity market scenarios they should test. There has been a general tendency to assume that some levelized growth exists in the fund for purposes of producing the testing. That tends to overstate the level of sufficiency.

My goal is to come up with a model and a modeling process that fairly represents the level of sufficiency. Why is that important? Say that if you're combining these results with other product lines, particularly general account product lines, and trying to form a conclusion about the adequacy reserve. If the margins of those other products are fairly narrow and have some exposure to interest rate risk, the combination might give you a false sense of security. The goal here is to produce variable results that, when combined with fixed results, will give you a greater degree of comfort in forming your opinion.

As it relates to the variable products that have guarantees in them, you want to bring in those risks accurately and have them reflected in the exercise. In the pricing context, it ultimately comes down to setting the price properly. Pricing variable products, as difficult as it is, focuses on how much capital you have, and what the measure of profit is. Trying to set prices in a static context also starts to lead you to conclusions that may not be entirely accurate, particularly as it relates to the period-to-period earnings pattern. So, if you build a model that reflects the true market and product dynamics, you'll get a better sense of the level of profit and increase the probability of meeting the profit objectives. If you're dealing with a guaranteed minimum death benefit or some of the other guaranteed benefits, you definitely need to reflect the inherent risks in such a way that if you ever decided to hedge those risks in the financial market, you can do so.

Let's say you want to price out a benefit that is 15 basis points. If you try and hedge it in the financial markets and find that it costs you 20, something's wrong. If you want to eliminate the risk, you have to take a loss. Finally, from a financial and risk management context, our recent experiences with the equity markets are giving new meaning to what one-year GAAP earnings volatility really means. In fact, we can even talk about monthly GAAP earnings volatility. So, if nothing else, building a good platform provides a basis for giving management some perspective on how bad it could get.

And the risk is not just near-term, but also long-term. Any kind of sustained bear market can have an interesting impact on the profit stream and stock price of a company that has a significant book of business in variable products. The ancillary guarantees have substantial tail risks, and I'm going to illustrate that a little later. The framework can help manage earnings, volatility, and risk. In one context, managing comes down to deciding whether we want to hedge our exposure. Do we want to protect earnings? Do we want to protect our risk exposure on some of the guaranteed benefits? Without that modeling platform, you can't answer those questions. In the real world equity markets go up and down, bond markets move up and down, and we can experience worse-case scenarios like a bear market and low interest rates. History tells us that the markets move in a correlated way. We also have some insight on policyholder behavior as it relates to lapse and transfers.

Many of you have noticed that we had all this volatility over the recent months, and the asset allocation in most mutual funds and variable products hasn't changed all that much. There has been some outflow, but nothing traumatic. There is stability, but building that in and understanding what could go wrong is important. It's part of reality. For the guaranteed benefits, as well as for your profit stream, mortality risk is not something that should be ignored. Another variable annuity reality is asset allocation—how the deposits are allocated among the different accounts. Remember, you're not dealing with the S&P 500, you're dealing with XYZ growth and income funds, and their performance is not going to be the same as the performance of the S&P. Odds are it's going to be worse, as most managers fail to outperform the markets.

A trend in the industry is that we tend to build simple static models. In doing so, we end up with a pricing exercise driven by how large our mortality and expense (M&E) fees are, our lapse assumptions, and our expense assumptions. As I said, we have a tendency to overstate adequacy and we end up with models that have little utility for financial management and no risk management utility. We, as a group, have generally recognized that when we start guaranteeing benefits, we need to get a little more sophisticated, so those models tend to be more sophisticated.

What are the attributes that go into an ideal variable model. First, we want a stochastic process along each major asset class. Imagine that you can have 50 funds. I'm willing to combine all those that are domestic equity into one. I'm willing to combine the international equity of the bond accounts, the balanced accounts, and so forth, to end up with asset-class-specific stochastic scenario generation that is allocation-distribution specific. Believe me, when I have 100% equity and 0% everything else, my response to a market correction is much different than when it's 50%/50%.

The model should also have dynamic lapse and transfer. You must factor in the fact that it's not the S&P that's growing or declining over time, but the actual funds in your portfolio.

Let's build the model. We include issue year, issue age, and asset allocation distribution cohorts. For every issue year and issue age, I'm suggesting that you separate the people who are 100%/0%, those who are 0%/100%, and those who are 50%/50% and treat them differently and analyze the different return dynamics. Somebody who is 100%/0% today could have been 0%/100% two weeks ago, so his or her return since issue is very different from that of the guy who was 100%/0% since day one. Table 1 illustrates what I'm trying to present conceptually. It shows the same issue year and issue age, but different allocations and returns on a year-to-date basis. You can arrive at the returns fairly easily. It's the current account value relative to net in's and out's. They're likely to exhibit different behavior. I think there's enough experience to suggest that the behavioral differences aren't going to be in small bands; they're going to be in the extremes. And those are the ones they're probably most concerned about.

TABLE 1
Liability Structure
Cell Structure

Issue Year	Issue Age	Equity	Fixed	Return
1994	55	100%	0%	32%
1994	55	0%	100%	6%
1994	55	50%	50%	20%
1994	55	50%	50%	8%

At some point, everyone might jump out of the equity market. Interest rates might be high enough that moving to a fixed option seems like a good idea. Surrender charges might end, and the person who sold the product to policyholders will convince them it's a good idea to go to another variable annuity. I'm not sure why they would want to do that, but I'm sure there's some salesperson who has convinced someone to do that.

Cumulative return affects behavior, how the assets are allocated, and returns.

Table 2 is a little complicated. It shows a variable product with equity and bond funds. This table identifies year-to-year returns for the equity account. There's an annualized return of 10.9% over that four-year period, and the same thing applies for the bond, but it's 6.7%. Take three policyholders: Dick, Jane, and Spot. Dick is 100% equity/0% bond, Jane is 70%/30%, and Spot is 30%/70%. When I compute their individual returns on a year-to-year basis, in aggregate, their return is 9.8% over the period.

TABLE 2
Liability Structure
Asset Allocation

	Year	**1999**	**2000**	**2001**	**2002**	**Annual**
	Equity	**20%**	**-10%**	**8%**	**30%**	**10.9%**
Policyholders	**Bond**	**6%**	**12%**	**7%**	**2%**	**6.7%**
Dick	100%/0%	20.0%	-10.0%	8.0%	30.0%	12.5%
Jane	70%/30%	15.8%	-3.4%	7.7%	21.6%	10.0%
Spot	30%/70%	10.2%	5.4%	7.3%	10.4%	8.3%
Total		15.3%	-2.9%	7.7%	20.4%	9.8%
Fees	Allocated	$41,500	$40,300	$43,400	$52,300	$177,500
	Constant Yield	$39,500	$43,400	$47,600	$52,300	$182,800

Notice the difference in their year-to-year returns. In some instances, it's very much different, and it's a reflection of their asset allocation. Then I calculated the fees that would be generated on an annualized basis, recognizing the specific return and asset allocation for one of the three policyholders, and produced the pattern shown. I took the 9.8% assumption, grew their aggregate balances forward at 9.8% on a year-to-year basis, and computed the fees. By assuming a levelized return, in aggregate, for these policyholders, the income pattern is going to be materially different: $2,000, $3,000, and $4,000.

Surprisingly, even the total fees generated during this four-year period are different. They're only equal in the last year because I'm doing the fees end of year, when all the values get together. There's a compound 9.8% annualized return in here, so it maxes out at the end. When you think about this in a stochastic context, you realize you can get some wild swings in terms of your fee levels because you're getting down to a more refined level of asset allocation, growth, and the particular account values over time.

I want to talk about disability structure transfers and their impact. In the end, one has to decide if the risk of transfers are really worth the effort, in terms of modeling. I have one policyholder with two accounts—account A and account B (see Table 3). The individual has $500 in each account.

TABLE 3
Liability Structure Transfers

Account	Amount	Period 1 Return	Period 2 Return	Period 2 AV	Period 2 Fees
A	$500	20%	-10%	$540	$5.4
B	$500	-10%	20%	$540	$5.4
				$1,080	$10.8
		Period 1 Transfer			
A	$500	$300		$810	$8.1
B	$500	($300)		$180	$1.8
				$990	$9.9

I make some assumptions about return over two time periods. Account A returns 20% during the first period, minus 10% in the next, and minus 10% and 20% for account B. So I bring the account value forward to the end of period two, and each account has $540 in it. On a 1% M&E charge, I generate 5.4 in fees and $10.8 in the aggregate.

Then we go to scenario two. In response to the returns in period one, we have a reallocation at the end of period one. Three hundred dollars is going to move from account B. The policyholder is going to bail out of that poor return into account A, then project forward to the end of period two. So I end up with $810 and $180 in the account value or $990 total, which is less than $1,080. I calculate my fees and, interestingly enough, I end up with less fee income.

This illustrates the point that, if you build in these dynamics to the extent that you can make realistic assumptions about how policyholders are going to behave in a transfer context, you're going to generate differences in the level of fees. This is a fairly extreme example designed to make my point.

Let's talk about scenario generation. It's important you build in a good interest rate generator. Depending on what you're doing and how you're doing it, it could be a single-factor model. In other instances, you might need a two-factor model. For example, if you're going to model fixed-account performance, you probably need a two-factor model. You need to reach out somewhere on the yield curve, and can't do that with a single-factor model tied to the short rate.

The scenario generator should be risk-neutral. If you want to do option pricing on those benefits or look at hedge strategies, that needs to be done on a risk-neutral context. An equilibrium context is for expected returns on the equity market of 12% with some volatility assumption, and so forth. Then, using that interest rate generator, you develop stochastically generated returns for each asset class that you built in, such as domestic equity, international equity, bonds, and balanced funds. That's done using mean and variance, expected return, variance of return for each of those asset classes, and correlation between each asset class and all the other asset classes. It results in each subaccount having its own return. If you then break your cells into asset allocations at the policyholder level, the cell-level returns will be different.

Table 4 is a standard correlation matrix that you would use to generate the different returns for the different asset classes shown here: money market, bond, balanced, and domestic equity. Each has its risk premium over the short rate, standard deviation, and correlation with the other funds. Then you can look up in any text how to use these correlations and the interest rate generated to develop the period-to-period returns on a pathwise basis for each of these funds.

TABLE 4
Scenario Generation

Asset Class	Mean Risk Premia	Standard Deviation	Interest Rates	Money Market	Bond	Balanced	Domestic Equity
Interest Rates		4.7%	1.00				
Money Market	0.0%	0.3%	-0.20	1.00			
Bond	1.6%	6.5%	-0.80	-0.05	1.00		
Balanced	3.5%	12.0%	-0.50	-0.02	0.70	1.00	
Domestic Equity	6.0%	18.0%	-0.30	-0.02	0.60	0.97	1.00

In the context of policyholder behavior, Morningstar data suggests that the equity allocation among funds over time has stayed in the 60%/40% and 70%/30% bands. That's not the individual policyholder level, but in aggregate, which would suggest that there isn't a lot of transfer activity. When equity markets grow 20%, that changes your allocation naturally. I don't know that people are saying, "I was 60%/40%, the market is up 20, and now I'm 70%/30%, so I need to reallocate back to 60%/40%." I don't think so.

What are the behavioral options in distress? Stay the course is one and there's some strong evidence to suggest that variable annuity policyholders do stay the course. 1035 to a new variable annuity? What are we going to do? Am I going to transfer my 70% equity/30% balanced fund allocation out of a variable account into a new variable account, get a new commission, a new surrender charge, and do 70%/30% again? I'm not sure I want to do that. Maybe there's a producer who can influence that outcome. Should I move to a fixed account? Maybe, but that, quite honestly, is a real risk. I may decide to move to a safe haven, but there is some age consideration. The average policyholder tends to be fairly old when he or she buys. Then policyholders get older and at some point die. So, if they bought in at age 65, 10 years later they're 75. I don't profess to know what their behavior is going to be 10 years later, but I think their mindset will be a little different than it was the day they bought the product. The more likely outcome, particularly under stress, is a reallocation. There's going to be a move to safety to fix or balance the account. We've seen a bit of that in the press, and some studies suggest that there has been some movement. Maybe there hasn't been enough stress in the equity market, but if somebody shows up and announces a bear market, and we experience it for 12 months, there could ultimately be a delayed reaction.

Lapse in transfer functions can transfer dynamics and not independent events. There's probably some point where lapse becomes a better option than transfer, and vice versa. If you define a lapse function, it should be a function of recent and historical account performance. If the market moved down 10%, and I'm up 40% since I bought in, what am I going to do? It might be somewhat different for the person who bought in and is breakeven since purchase. So the current shock, whether it's up or down, and where I am since issue probably is going to influence my behavior.

How my assets are allocated are important. If the market is down 30%, and I'm 100% equity, I'm going to think about it. If the market is down 30%, and I'm 50%/50%, maybe I'll think about it. If I'm 100% thick, I'm a smart guy. There are fixed-option alternatives. If you're in a variable annuity, the fixed option looks awful. But, if somebody out there is waving an interest rate that's 1% higher, that will drive the dynamics, and that scenario could occur. And the distribution system always requires some consideration. If you're thinking about building in lapse and transfer dynamics in any variable annuity modeling, you need to give it some careful thought. If you do build it in, it at least gives you the ability to test sensitivities to different lapse and transfer dynamics. So, in summary, my ideal variable annuity model incorporates stochastic correlated scenarios, fund-specific or asset-class-specific growth, and allocation of distributions of policy at the cell level. It permits the projection of income and benefits in a more realistic way. One can always argue that there is more complexity, and less ability to interpret the results, and I recognize that.

I want to talk about modeling as it relates to guaranteed benefits, using a guaranteed minimum death benefit as an example. A risk on this benefit depends on a lot of elements: how you design the product, what you expect to happen in terms of mortality, what you expect the policyholders to do, how the assets are allocated, fund-specific performance, transfer, and lapse.

Chart 1 is an illustration of the risk profile for a guaranteed minimum death benefit. I'm not going to spend too much time describing the benefit features in the assumptions, but it incorporates most of the elements I've been talking about. This is a present value of earnings illustration with the results across 400 scenarios rank ordered from highest to lowest. We've looked at three benefits. The bottom line is the return of premium, the middle line is the 5% rollup, and the top line is an annual ratchet. The cross-over point in about 60 of the 400 scenarios produces a mid-economic loss. It's hard to illustrate, but if you go to the tail, those worse-case scenarios look pretty ugly. But they're not impossible or improbable.

Our friends in Japan have had an experience that is worth noting. Arguably, it can happen here. Table 5 is the Nikkei index for 1987 through 1997 showing the year-to-year return. It had a big run up, followed by almost a 40% decline in 1990, and it has been taking hits ever since. This is a real world scenario because it happened.

On a cumulative basis, if you started in 1987, you're negative 23%. If you started in 1990 you're a negative 63%. These scenarios fall outside the tail of that distribution in Chart 1. If we model those scenarios on a guaranteed minimum death benefit, they would fall off the chart. My point in all of this is that these aren't risk-free benefits. They present a fair amount of risk, and you should probably build models that will give you a good sense of how much risk is there.

Table 6 is a simple illustration. These numbers are fudged, because I don't want to show you the real numbers, but the order of magnitude is OK.

We went risk-neutral with our assumptions, our cell structure, and so forth, and then option priced the benefits. The annual ratchet at 20% market volatility would have a benefit cost of about 25 basis points of account value if you only assume base lapses and no transfers. And surprise, surprise, if you assume base lapses plus a dynamic lapse in transfer function, it helps, not hurts. If you reduce the volatility from 25% to 20%, you get a significant reduction in the premium.

CHART 1
GMDB Risk Profile

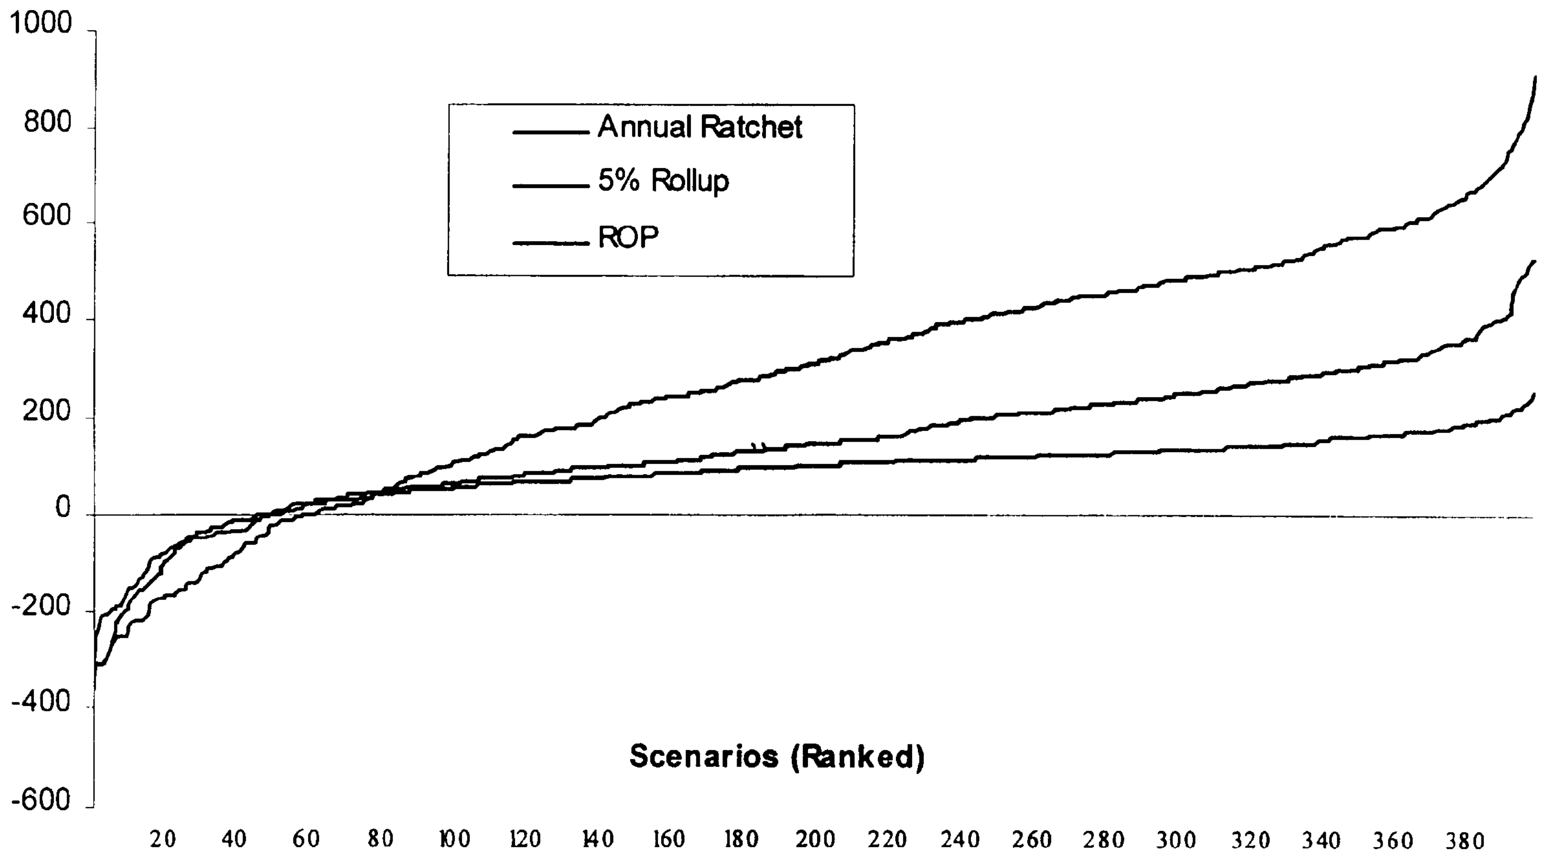

TABLE 5
Nikkei Returns

Year	Annual Return	1987 Cumulative	1990 Cumulative
1987	15.3%	15.3%	
1988	39.9	61.3	
1989	29.1	108.2	
1990	-38.7	27.7	-38.7%
1991	-3.6	23.1	-40.9
1992	-26.4	-9.4	-56.5
1993	-2.9	-12.1	-57.8
1994	13.2	-0.4	-52.2
1995	0.7	0.2	-51.9
1996	-2.6	-2.4	-53.1
1997	-21.2	-23.1	-63.1

TABLE 6
Option Pricing GMDB

Product	Market Volatility	Base Only	Base + Dynamic
Annual	25%	25	21
Ratchet	20%	18	15
5% Roll Up	25%	30	23
	20%	27	20
Return of Premium	25%	11	8

Let's say you had an annual ratchet priced at 18% and 25% of the market, and you wanted to hedge out the risk. After buying all your derivatives, your net cost is likely to be negative. With a 5% roll up on 30 basis points, there's a base lapse of 23 and dynamic assumptions, and it's less sensitive to the volatility. The annual ratchet is reset every time the market moves up and down, so you're in the money. The more volatile the market, the more likely you are to be in the money. With a 5% roll up, there is long-term growth in fund performance relative to that 5% accumulation assumption. The return on premium is 11% and 8%. I just like to illustrate the point that it can be option priced. You can look at derivative strategies to hedge out the risk, and if you approach the market from a pricing point of view, you want to have enough premium to be able to buy the derivatives you want to hedge. You can hedge the tails and buy a set of derivatives to reduce your tail exposure, but there's some cost to that. It brings down the right-hand side of that risk profile. It can also do things like dynamic hedging, where you're calculating the Greeks which are equivalent to duration and convexity in the derivatives markets. Finally, use those to identify and manage your hedge position on a going-forward basis.

To summarize, there is enough risk in variable annuities and their ancillary products to make serious model building worthwhile. Remember, if the model is poorly designed and constructed, it will provide little useful information. If it's well-designed and constructed, it will give you a wealth of valuable information.

MS. MEREDITH A. RATAJCZAK: I'll explain the modeling process by using some case studies of real-life experiences to give you a sense of the issues. I will talk about model structure from the standpoint of how many cells you need, what the issues are, and what you should consider. I will also cover validation techniques. Once you have your model built, how do you validate it, and the assumptions you have used? There are some realistic assumption selections for fixed annuity products, but I can't give you a cookbook approach to building an annuity model. The reason is that every model you or I build will have a different purpose. You might be under different time constraints for building and using the model.

Model building is more of an art than science. Constraints must be considered to determine how big your model can be, how quickly you have to use it, and what sort of data you need to back up that model. Our computers are very fast today. When I look at my annuity models today compared to those done seven or eight years ago, I do a lot less modeling now. It's very easy to get an in-force file or tape from somebody and, essentially, have model plans to fit every single model plan in that file. In some cases, that's not realistic, but modeling is a lot easier now, because you're not actually modeling.

Model size depends a lot on what you're going to use it for. If you're doing a quick-and-dirty estimate of value for a block of business, you can do that with a one-cell model based on publicly available data. If you're doing cash-flow testing, or even if you're coming up with closed-block funding estimates, you might have a very granular model with lots of cells to capture as closely as possible the underlying characteristics of the model that you're working with.

When I start any project that involves building a model for annuity products or life products, the first thing I do is determine what the model will be used for. If you're going to use it for that quick-and-dirty projection, you probably don't need to look into all the different model plans. Know about your product, in general, and come up with an average cell to represent it. Cash-flow testing requires a lot more cells, so you must consider your purpose.

Identify your constraints in terms of the data available to help you build the model, time limitations, and resources. Are you building the model, or will somebody else help you build it? Then you gather model information. To determine what your model cells are, first get some statistics on the cohort of policyholders that you're modeling.

Typically, I get a seriatim in-force listing from a client; then get my students to set up an Access database by model and plan code. The reports tell me the distribution of the business by issue age categories; it is usually 10-year age bands and issue-year distribution. I use those statistics to determine my major plans, issue ages, and durations.

Typically, I don't collapse across duration. Duration is a very important aspect of a deferred annuity model. Even if there's only a little business in it, because computers are so fast, I keep the durational perspective in the models. Then, I review the product from the standpoint of gathering information about surrender charge patterns. Are there any special features in the products, such as a nursing home rider? You have to look at the type of information you're going to need to support the model. Then I select major plans based on the statistical information the students gathered. For a typical modeling situation, I choose plans that represent about 80–90% of the in-force business. The other 10–20% is usually scattered in several small plans that might have characteristics similar to the major plan.

Then I map those minor plans into the major plans based on similar characteristics and build the model. It sounds easy, but usually it's very time-consuming. In a typical cash-flow testing project, model building can be a very time consuming aspect of the whole project, especially if you're not familiar with the block of business to begin with. The project I'm working on has hundreds of plan codes that I need to understand in order to build the model. It's very important that my model has the right amount of granularity in it, or the right number of cells, so I can manage the process. I have to go through the plan code listing and the product encyclopedia and figure out the model cells.

As far as the model structure itself, the key characteristics that you should reflect in your model building are plan and product characteristics. If you have a flex annuity and a single premium deferred annuity, you don't necessarily want to collapse those into one plan. They have different premium characteristics, different surrender charge structures, and probably different lapse characteristics.

Some companies have many alternative distribution channels, so I use the distribution system as a model characteristic. If you have products that are sold by a career agency system, and others sold in a brokerage environment, things like commissions, lapses, and expenses will differ. Some kind of a distribution channel indicator is typically found in the valuation and the administrative system extract. For deferred annuities, mortality is not a key variable in terms of its impact on the results.

You will find you're modeling older issue ages. My deferred annuity model typically has cells with issue ages of 45, 55, 65, and sometimes 75. You don't often see many younger ages, so I bucket those as the youngest cell. We are doing age differentiations because people are adding some new riders such as the nursing home rider, where age does make a difference in terms of utilizing the rider in your assumptions.

With one deferred annuity model I had on my computer, I started off with the in-force file. I want to show the impact on present value of profits (PVP) if I changed some of the characteristics.

- All cells collapsed into one average cell – PVP @ 10% = $136,352
- One average issue age with duration distinction – PVP @ 10% = $129,310
- Issue age and duration distinction – PVP @ 10% = $113,676
- Issue age, duration and distribution system distinction – PVP @ 10% = $120,505

I took my in-force business of $136,352 and I came up with an average issue age, an average duration assumption, and because there are distribution channel differences, weighted assumptions for lapses and expenses.

Then, using that same average issue age, I put in the duration distinctions. In this particular model, there were three durations, 1995, 1996, and 1997. In this particular level interest scenario, I'm still picking up some dynamic lapses, so reflecting the durational differences did make my PVP go down.

Next I have issue age and duration distinction, so I'm not collapsing. I had ages 35, 45, and 55, and I'm reflecting all of those. Finally, I've separately modeled the distribution differences in the model. There are not huge differences in the values. But let's say you go from something where you have shown none of these characteristics and distinctions in your model, $136,352, to where you're showing issue age differences and duration distinctions down to $113,676. If you're using your deferred annuities to offset your immediate annuities, say, in cash-flow testing, it's possible that you could see swings in your results from positive to negative, depending on how you set up the model You want to capture the distinctions that you have in your portfolio realistically.

Once you've built the model, you need to validate a number of things. Imagine that you have a model that has 10,000 policies in it, an account value of $1 million, and reserves of $1.2 million. When you build this model and push it through your projection system the first time, first check to make sure that you have modeled 10,000 contracts. If your projection system calculates a starting account value in reserve, you want to make sure that your model has a good fit.

You should do an actual-to-expected analysis to make sure that your model is an appropriate representation. To determine your tolerance level, you might want your starting account value to be within 99% of the actual value. Tolerance levels will change based on your purpose. If you're doing the quick-and-dirty analysis with one cell, your model fit might not be as good. But for cash-flow testing, you will expect to see a fit in the 98%, 99%, or 100% range.

The next area of validation is your assumptions. After I've built the model and determined my assumptions, the first thing I would do is run my projections for a year or two and see what income statement items look like for this particular model. If you are projecting premiums, you might look at how realistic the calculated premiums look compared with your 1998 annual statement. If you have $25 million of premium for 1998, and your projection shows $15 million, I would suspect that something is wrong with the model.

I also try to validate the fund development mechanics. I look at the relationship each year between items such as reserve-to-account value and cash value to account value to see if the implied surrender charges make sense. To validate dynamic assumption mechanics, make sure you have a dynamic formula that is calculating what you think it should be calculating. That's your first level of validation. For the next level of validation, let's say that from December 31, 1997 to December 31, 1998, the interest environment is down 1%. You can go back to last year's cash-flow testing and calculate a down-1%-over-ten-years scenario. Based on last year's projection, you look at what type of lapses the model is producing based on those assumptions. If we've been in a down-1% environment, and you can show actual lapse experience for this particular block of business, and how closely your dynamic assumptions are replicating reality. It's complicated and squishy, but that's one thing that you can do in terms of validating your dynamic assumptions.

The more difficult validation, as far as dynamic assumptions go, is to show how closely these formulas replicate management actions. There have been situations in a new consulting assignment where we ask people about their crediting strategy and how management reacts in different interest rate environments. Those are difficult questions for some people to answer. They haven't thought about how their business operates in those terms, so this level of validation is very difficult. In any of these situations, look at the realities of actual assumptions in relation to how you have defined them based on a formula.

What are the key considerations for making realistic assumptions? First, look at the products. If you have a product with a bail-out feature in it, you want to capture the fact that, if you set the credited rate lower than the bail-out rate, you probably will have a lot of lapses. Then look at how your policyholders might behave. In today's environment, interest rates are very low and there are not many alternative vehicles, so lapses might be low. But think about how your policyholders might react if you reduce your credited rate or if the interest environment goes up or down.

You also should consider how management behaves. That's a difficult one. We can't always get to that easily. And this whole assumption development process should result in realistic assumptions.

I'll talk a bit about investment strategy. If you're doing cash-flow testing using an investment strategy that doesn't closely mirror what you're doing now, it's not realistic. Realistically assess your credited rates, market rates, and lapses for your business. Fixed deferred annuities aren't as exciting as what Frank was talking about. The key assumptions are defining your market rate. Most people have lapse formulas that have a market rate component in them. How do you define them? I've seen companies define them in many different ways. Renewal crediting strategy is also important. If you do sensitivity testing, you will get variability in the results.

Companies usually view the market rate as who their competitor is. That may be defined in terms of what's going on in the new money environment. The market rate might be defined as the seven-year Treasury plus a spread. And that's what the company defines or views as its market rate. It might also be based on some sort of a portfolio rate, or a company's competitor might have a variable product, because where's the money going to go if it doesn't stay in the fixed product. So the definition might be some sort of a fund growth assumption. There are advantages and disadvantages to all of these, and you might find companies that blend these approaches to defining how they view the market. When you set up your formula for defining your renewal rate for cash-flow testing, set crediting strategies that replicate the past exactly. That isn't easy to do because many factors enter into a company's decision for setting renewal interest rates. Part of your renewal rate formula might have something to do with how management is looking at the market. Perhaps it wants to subsidize the renewal rate a bit. Perhaps it wants to give high first-year rates and low renewal rates. In these cases, you try to replicate, as best as you can, what actual experience has been in different interest environments.

You have to capture the company's philosophy. Historically, if you know that interest rates have gone down 1% in a year, once again, you can track back to last year to see what sort of credited rate you might have projected to be calculated based on the company's portfolio—or however it defines the renewal rate—and see how it compares to where the company actually is today. That's one way to do a reality check on your assumption.

I've seen and used many different formulas for setting the renewal credited rate. You can link it to market movements. Some people set their renewal rates as a market rate less some spread, trail the market up, and immediately follow it down. Some people say it's 175 basis points off their portfolio rate. So there are many ways to set renewal rates. It can be linked to a market rate, to the company's portfolio rate, or to some sort of combination of those.

In terms of constraints, you might have the guarantee rate to take into consideration. Companies might say they are not going to reduce their credited rates more than x% in a year, so that's factored into the formula. They might say they won't let it go any higher than a certain number based on the scenarios that you're testing.

Once again, lapses should be consistent with the actual experience. The way to do that is by conducting a persistency study to see what the lapses are in relation to what you might be projecting with your formulas. That's a relatively straightforward exercise. It should consider the policy's characteristics. Let's say you have a cliff surrender charge; it's 6% for a while, and then zero. I would suspect that at the end of that surrender charge period you're going to see some pretty hefty lapses. Experience has shown that if you have bail-out features in your fixed deferred annuities, you probably want to have a component in your lapse formula that says, "When you get to the end of the bail-out period, if they do set the credited rate below the lapse rate and trigger that bail-out provision, you can expect to see increased activity in lapses." So, the policy characteristics need to be considered in the lapse formula. The lapse rate should be as consistent or realistic as possible.

We've used many different lapse formulas. The one we use today has a market rate/credited rate component that might have a scaling factor on it and an exponent. It also has a reduction for surrender charges, because experience has shown that the incidence of having a surrender charge in-force keeps policyholder lapses down. Durational differentiation of lapse assumption is necessary. If you've collapsed your model into one or two durations, and have things like surrender charge adjustments in there, you will not appropriately capture the actual lapses for the in-force business.

I've taken a block of business and modified the lapse formula. The lapse behavior should consider how companies set credited rates. If credited rates go far below market rates, you can expect a lot of lapses associated with that. For deferred annuities, I've shown what happens if you assume investment in seven-year noncallable bonds or mortgages. Results for an exercise in which you're modeling assets can be very sensitive to changes in the underlying investment and reinvestment assumptions.

I'm not going to go into any detail on that. I just wanted to point out that when you do modeling, and are looking at the asset side of the house to match the liability side, you must examine how the company is actually investing. If the company is buying mortgages and collateralized mortgage obligations (CMOs), and some equities, and you're only modeling bonds, how realistic are your results going to be if you're not actually reflecting the interaction between the assets and the liabilities? In the past, when people were using systems that were less sophisticated in terms of asset modeling capabilities, you might say they only represent 5% of what we have, so we're going to make some simplifying assumptions. I'm not saying that that's a bad thing to do. Most of us still do that sometimes in the interest of time, but I think we've gotten a lot more sophisticated, both in terms of the tools that we have available and the way in which our investment departments are investing. If we polled everybody, you'd find that people are spending more time coming up with very detailed models on the asset side to go along with their liability models.

The assets you are modeling should reflect what your investment department is actually buying. I'm sure the investment people tell you they're buying seven-year bonds that have a particular spread and are of a particular quality. Those are important facts to keep in mind whenever you're setting up your reinvestment and disinvestment assumptions. I've seen all kinds of strategies—buy bonds, mortgages, some combination, CMOs. People have options, so you see many different investment and reinvestment strategies. Systems today allow for a lot of sophistication in terms of what you buy and how you sell. In terms of disinvestment, you might see a pro rata sales strategy, or one that maximizes capital gains and minimizes capital losses. In the system I'm most familiar with, you can set priorities for what assets to sell first. Or, if I want to take into consideration the fact that I can only have 20% of my assets in noninvestment grade, I'll sell these first or buy something else. But some people still assume a buy-and-hold strategy and, if necessary, employ a borrowing strategy instead of liquidating to cover cash needs. The problem with that is, in a situation where you need to sell year after year, that short-term borrowing becomes long-term borrowing, and you might not be assuming the right loan rate.

What happens if I make a modification to my lapse formula? In one model, I had two issue ages, three issue years, and two distribution channels. Expenses are $55 per policy for channel 1, and $40 for channel 2. One might be career agent, and one might be a brokerage. Channel 1's lapse rates start at 1.5% and grade up to 10% for the year right before the surrender charge disappears. They fly up to 30%, and then stay at 15% thereafter. That's the base lapse assumption. For channel 2, the base lapse rates start at 3%, go up to 15%, hit 35%, and then, ultimately, stay at 20%. This might mimic what's happened in the brokerage community. Credited rates equal the portfolio rate less 175 basis points—a pretty simple renewal crediting strategy. You are assuming 175 today. You may want to drop that in the future if new money rates continue the way they are, but we're using 175 here. The interest-sensitive lapse formula that we used has a market rate minus a credited rate component with a scaling factor and an exponent. And we do a reduction based on the surrender charge divided by four.

The asset side is simple. Just have some callable and noncallable bonds. We sell assets pro rata. We are purchasing bonds and mortgages because this company has changed its investment philosophy. So we've assumed we're earning 50 and 85 basis points over Treasury on those securities.

I took this simple model and did the New York Seven or the NAIC Seven Scenarios on it based on an interest environment a little steeper than it is today (see Table 7). For this particular block of business, the pop-up 3% is the bad scenario. If you're aggregating, you have some pretty good margins in the other scenarios. As far as lapse formula adjustment, there is an impact on the result from getting rid of the surrender charge piece. For the down scenarios, where you're floored at a minimum lapse rate, it doesn't have any impact. For the up scenarios, if you look at the up/down going from 328 to 217, that's a significant difference. If you don't include something like that, and you are using your deferred annuity results to help boost your immediate annuity results, then you might be missing out on some surpluses you have available for that offset.

TABLE 7
Assumption Development Case Study

	Base	Lapse Formula Adjustment
Level	1,123	1,088
Up	836	786
Up/Down	328	217
Pop-Up	(400)	(564)
Falling	1,135	1,133
Down/Up	1,028	1,028
Pop-Down	1,633	1,633

The important thing to remember is there is no cookbook approach to building a model. You have to consider the purpose for your model, go through the actual model development process, and validate not only your model, but also the assumptions you're building. The key for fixed deferred annuities is developing your assumptions in terms of a formula. Try to be as realistic as possible in terms of how you are modeling policyholder behavior, management behavior, and the underlying mechanics of the product.