RECORD, Volume 24, No. 1*

Maui I Spring Meeting
June 15–17, 1998

Session 75OF
Bad Behavior: Policyholder and Interest Rates

Track: Investment/Product Development
Key words: Investment, Finance, Forecasting, Product Development

Moderator: KLAUS O. SHIGLEY
Panelists: KENNETH P. MUNGAN
Recorder: KLAUS O. SHIGLEY

Summary: Many interest-sensitive life and annuity products have embedded options which value is a function of policyholder and interest rate behavior. While some data exist that help project their behavior, there still remains a tremendous lack of data in rising rate scenarios that might help determine the value of those various options. Despite that, much work has been done to better understand these embedded risks and determine how best to deal with them.

Mr. Klaus O. Shigley: Ken Mungan was formerly with Allstate where he developed some extensive experience modeling customer behavior for insurance products. Ken received his B.S. degree from MIT. Both Ken and I will cover annuity products. I'm going to follow Ken with a presentation on participant behavior, the participant behavior assumption in the stable value arena and how we quantify the option risk in that business.

Mr. Kenneth P. Mungan: I am from the Chicago office of Milliman & Robertson (M&R). I'd like to speak about policyholder behavior, by focusing on annuities and universal life. First, I'm going to talk about the systematic themes by which people collect and utilize information on policyholder behavior. Based on what I see companies do with policyholder behavior and lapse studies, I would say many people are following the herd. They simply take whatever formula's in their financial models straight out of the box and maybe jigger the coefficients around a little bit, or maybe not, and then they're good to go. They go on and use that in their financial modeling without any adjustments at that point. I also think many

---

Note: The charts referred to in the text can be found at the end of the manuscript.

companies might not have systematic processes for gathering information on customer behavior, so I'm going to talk about the need to work on those two points.

How many people feel that they might fall into this category, that they're kind of going along with traditional formulas without systematically collecting and using information on lapses? Good, I see a couple of hands going up. What I'd like to do is go over some things that would be a little bit more bold and adventurous when it comes to studying lapses so you can jump out and do something new and launch your company into a new era of understanding customer behavior. This is very important, because if you have a better understanding of liabilities that means you can develop better investment strategies that will give you more stable and secure financial performance.

Before we get into the actual mathematics of modeling lapses, it's important to understand that lapses and customer behavior, in general, are part of a dynamic process. This is a process that you're probably all very familiar with. You may have a marketing strategy and that marketing strategy is going to feed into a dynamic system that involves the crediting strategy, lapses, and investment strategy all interacting with one another. We're all very familiar with and focused on the financial results that flow out of that system.

In addition, information flows out of that system. There's information on the customers' behavior and then the loop starts all over again. If anyone has been involved with a quarterly close process, he or she knows that this loop goes on, and on, and on. The question is, Are you just walking through this loop time after time or are you systematically collecting and using the information that comes out of this process?

To do that systematic collection I recommend a lapse analysis process overview. Let's start here with making predictions. Let's say you already have some models and a new quarter or a new year is starting and you make some predictions to see what you think lapses might do. This prediction could take the form of a lapse function or it could be some predictions along different interest rate scenarios. Then you should collect information and update the database. You could update the database and get information on which policyholders lapsed, and a series of variables that you might want to use to predict those lapses.

Once you have all those data, you want to evaluate those predictions and, after evaluating, explore strategic implications. You could have more or less sensitive policyholders than you imagined, and that could have implications for your investment strategy, for your crediting strategy, and so on. At that point, you want to make modeling improvements. What I'm speaking of when I say modeling is mathematical modeling. You need to update your mathematical lapse formulas to reflect the new information that you gained in this process.

The final step in the lapse analysis process is implementation into a financial model. ALFA is M&R's model, but it could be whatever model that you're using to generate financial information over a large number of scenarios. Then you collect and analyze that information, and the process starts again. This process could go around every three months, six months, or one year. You have a discipline and rigorous process for collecting and analyzing information, and without it, you won't have the best investment strategies and your whole system of reacting to a crisis, such as an interest rate spike or a severe interest rate drop, probably won't perform in the best manner.

Does anyone in the audience believe that they have a system like this? Does anyone have a process for routinely collecting and using this information? We have one hand. Does anyone feel that they come up short in this area? There are many hands going up. Before we even get into the mathematics of modeling customer behavior, it's important to understand that you can have the fanciest formula in the world, but if you don't have a rigorous and disciplined process for collecting and using information, you could be fooling yourself. You could be just guessing how interest sensitive your business is.

Let's talk about a general framework for understanding customer behavior. To come up with that framework I've developed three different factors which influence customer behavior that you can use on any of a number of products; random, demographic, and strategic factors. The random factor for an annuity block could be interest rates. If you're selling direct response products, it could be things like the unemployment rate and inflation rate. Demographic factors could reflect the marketing strategy that you're targeting or the make up of the block of business that you might purchase from another company. Age and wealth level would be very important for annuities. Life cycle and gender could be very important for direct response type products. Finally, what variables reflect the strategy that your company's following? What market are you targeting and what crediting or pricing strategy are you following out? And then for a direct response example are you using telemarketing or direct mail and what is your billing method?

Once you've walked through for the product that you're responsible for managing what are the random demographic and strategic factors? The next step is to build a comprehensive financial model. You're going to have many scenarios that model

the random factors and then a financial model that reflects the demographic make up of the block of business that you're responsible for, and the current strategy of the company.
An audience member brought up that the equity markets have a very strong influence on lapses and shouldn't that be reflected? It definitely should. If you are using a financial model that generates interest rate scenarios and equity scenarios, concurrently, then you can do that in a very straightforward manner. I think a lot of people probably aren't doing that though. And if they're using just interest rate scenarios, then what you can do is have a driver of lapses that is not related to the annuity market rate, and that could be used to reflect equities as you mentioned. It is something that's pulling the money out of the products that isn't interest rate related.

Chart 1 shows lapse rates that come out of the lapse functions that I've seen. Policy years are shown along the x-axis and we can see the envelope of lapse rates that you would see over hundreds of interest rate scenarios. Let's say you ran 500 interest rate scenarios and for year seven, because that's the year in this example where the product had no surrender charges, there are shock lapses. At the 90th percentile level, lapse rates got as high as 50%. Those would be in the up interest rate scenarios. At the 10 percentile level we have low lapse rates, around 20%, and that's what I'm calling the dynamic formula.

In constructing this dynamic formula my goal was to get the correct year-by-year estimate of the mean lapse rate, and the right dispersion for each year about that mean. You want to get this magnitude and timing of lapses correct. I have, for comparison, a traditional formula which is an alternate formula that I've seen that had very little interest-sensitivity in it. You can see the 90th and 10th percentiles are much closer together. If you're using a formula that has very little variability in terms of lapses, you could be underestimating the actual impact that's going to have when you get into an interest rate spike situation.

Chart 1 is useful for whatever financial model you're using to get an eyeball estimate of the interest rate sensitivity that's embedded in your lapse formula. If you look at this chart, and you see the 90th and 10th percentile lapse rates right on top of one another, you might want to think about it a little bit to decide if you might have some more sensitivity.

From the Floor: I have a question on that chart. I see you have the10th and 90th percentile traditional lapse rate but they are still distinct. What is it that creates the distinction? Is it a different factor altogether?

Mr. Mungan: These are two different formulas. They're both reflecting the interest-sensitive lapses, but the traditional has very little interest sensitivity, so the coefficients in the formula would be dampened relative to the dynamic.

Table 1 is an example of the interest rate sensitivity you have on a particular product in a particular distribution channel. The goal is to have consistency of lapse modeling across all your products and distribution channels. For any of you that work for fairly large companies that have multiple distribution channels and a large number of products, it can be quite a challenge to get consistency. And if you have different groups of people that don't always communicate in the best way with one another, you can have completely different estimates of what the interest sensitivity is floating around your company.

TABLE 1
CONSISTENT LAPSE MODELING
ACROSS PRODUCTS AND DISTRIBUTION CHANNELS

Steady State Methodology for Examining Consistency
Issue $100M of new business each year for 11 years
Determine the steady state annual lapse rates on the entire book of business.

Distribution Channel A, Product 1
Scenario—Credit the market rate each year

Average Account Value $40,000

Years Completed	Credited Rate	Market Rate	Annual Lapse Probabilities	Survival Rate	Surrender Charge
0	6%	6%	3%	97%	8%
1	5	6	4	93	8
2	5	6	5	88	7
3	5	6	5	84	6
4	5	6	5	80	5
5	5	6	5	76	4
6	5	6	40	45	0
7	5	6	25	34	0
8	5	6	12	29	0
9	5	6	12	25	0
10	5	6	12	22	0
Steady State Annual Lapse Rate				9.0	

Table 2 describes the steady state methodology for consistency where you would issue $100 million of annuities each year. This example issued that for 11 years. You want a fairly long time period so that, at the end of the time period, the business remaining is small. The goal is to determine the steady state annual lapse rate on the entire block of business at one point in time. At any point in time, you've been issuing $100 million every year, and you have lapses associated with each of those cohorts at the current time. The table shows years completed, the credited rate today and the current market rate. Each of those cohorts is going to have an annual lapse probability that would come out of your lapse function, and you can calculate what the survival rate is for that cohort. Then you come up with a total lapse rate for the whole block of business, which I call the steady state annual lapse rate.

You should do this calculation for several different scenarios. You could do the calculation that I just did for a year, at 9%, where we're crediting 100 basis points below the market. That was distribution channel A, product one.

TABLE 2
COMPARISON OF STEADY STATE ANNUAL LAPSE RATES

Distribution Channel	**A**	**B**	**C**	**D**
Product	1	2	3	4
Crediting Scenario				
Credit the Market Rate	6%	5%	7%	6%
Credit 100bp Below Market	9	8	12	9
Credit 200bp Below Market	18	16	20	17

Differences should be explainable by product features and distribution channel characteristics

Then I ran three scenarios where we vary the credited rate. You can see a low, steady state annual lapse rate of 6% when we credit the market rate, 100 basis points below and 200 basis points below market. Then I've done the calculation for all the different distribution channels and products that the company has. Table 2 shows there are some differences here and that's perfectly fine. You don't necessarily want to have the same lapse formula throughout the entire organization, but the differences that do show up should be explained by different distribution channel and product characteristics. For large differences that you don't feel you can readily explain, you want to go back and question the assumptions that you're using. This is all about having a consistent process for collecting and analyzing information to make sure that you have consistency.

Notice, that we haven't even gotten into mathematics yet. Let me throw out one more question for the group. Are there any actuaries who feel they have several different distribution channels, many products that have been faced with this problem of having consistency? Does anybody want to comment on how they handle that?

Mr. Ronald L. Ziegler: We haven't tried to look at what you've termed as steady state, but we do annually look at actual lapses between two large distribution channels versus expected lapses under the given interest rate scenarios we actually lived through. We've done that for a number of years and you start to validate the model's reality at the time.

Mr. Mungan: The processes that I described are very beneficial. If you wait until there's a crisis, you've really missed the opportunity to progressively get more and more confidence with your model.

Let's go onto some math. I'm going to talk about two different ways of modeling lapses. The first one is something that you've probably seen, at least some variation of it. This is a traditional actuarial lapse formula for coming up with lapse rates. It's widely understood. A version of this is presented as an example in New York Regulation 126. Let me go through some of the variables. You have a market rate (MR), a credited rated (CR), some surrender charge (SC), and then some interest sensitivity factors, which I call Z (multiplier) and K (exponent). There is also some factor, A, that determines how your policyholders are amortizing their surrender charges. This type of formula would develop excess lapses, and then you could add in base lapses separately.

The amount of excess lapse =

Z * [(MR – CR) – A * SC] ^ K	if MR> CR + A * SC
0	if CR < MR < CR + A * SC
–Z * [MR – CR] ^ K	if MR < CR

The above formula might be familiar to you. You have some multiplier times a surrender charge, and adjusted gap between the market rate and the credited rate raised to a power. If your credited rate is substantially below market, that's going to give you a positive excess lapse. Maybe you could have zero lapses if your credited rate is within a certain band of the market. Then there are negative excess lapses if you're crediting substantially above the market. I've also seen some people zero out that last component. They assume there's never any negative lapse.

I've left this pretty vague in terms of all these variables. I haven't put in specific numbers, because you would need to come up with numbers for the Z, the K, and the A, the interest sensitivity factors, and the amortization of a surrender charge that's appropriate for your company. I have kind of a trick question, so we'll see if anybody wants to volunteer and help me out with this. Let's say I give any actuary in the audience a database. The database contains information on one million policyholders over the last 15 years and in all kinds of interest rate environments, reflecting many different crediting strategies that your company had pursued over that time. You have all of these variables—the market rate, the credited rate, the surrender charge on the policies. How would you go about determining the correct Z, K, and A factors that give you the best fitting lapse rate function to use in your financial projections? Has anyone had to do this, and how do you go about fitting parameters to a formula like this?

From the Floor: Regression.

Mr. Mungan: Regression? Yes, and it's hard to use regression. When you come up with formulas like this it can be very hard to use regression. That's the challenge that a lot of people have had to face. Certainly, it's not linear regression. Then you can go and make some transformations, but then you're stuck with all those if statement on the right-hand side. It can get pretty messy, but you've got to come up with something. You work with the database and try and use some kind of progression, but it's not clear cut and I haven't seen a lot of consistency from the different people that have had to handle this.

I'm going to go into a completely different type of lapse formula and into one where it's much easier to match up the data with the best fitting lapse formula.

Mr. Gregory William Chicares: Just a comment on that formula. I understand it would be hard to solve that with typical regression techniques. When I see something like that, with just a few parameters, the thing that occurs to me would be simulated annealing. I wonder if you or anyone here has tried that approach.

Mr. Mungan: I've used genetic algorithms, which are very similar. Whenever you have an ugly formula like this, you can use any genetic algorithm program, and it will rapidly converge to the best estimate. You've got to supply the program with how to determine the best estimate and you could do that using traditional statistical measures like the squared difference between the actual and expected lapse rates. I could talk for two hours on those techniques.

Let's go on to discuss logistic lapse models. This is something that most actuaries are unfamiliar with, but they're extremely useful. Logistic models apply an S-shaped curve to binary data, and all binary data are on/off data. It's a widely used statistical technique. Let me give you an example. Let's say I said you are actually here for a study on the effect of diet and exercise on heart disease. I'm going to study you for the next five years, and each year I'm going to collect information on an on/off variable. You either had heart disease or you didn't have heart disease, it's a one or zero. Then I'm going to collect information on the type of food you eat and the amount of exercise you get. Those are continuous variables and they can be binary as well. At the end of the study, I'm going to come out with a model to predict the probability that you're going to get heart disease based on the types of food you eat and the amount of exercise you get. Many of you probably have heard of lots of studies like this. Everyone has an understanding of what the results of that study would be. These techniques have been widely applied to the field of medical research where you have exactly that set-up. In each time period people either come up with some ailment or they don't. The researchers are looking for factors that might predict the group of people that would be most likely to get the illness.

What I've worked with is ways of taking these statistical techniques that were developed for other fields and applying them to the lapse problem for the insurance industry. In that problem, for each time period, the policyholder either lapses or they don't and the predictor variables would be things like market rate, credited rate, surrender charge, and the type of crediting strategy you're following or maybe the account value. You can put any number of predictors in there and then the model will predict the total lapse rate base, plus excess.

I'm going to walk through a simple version, which is a three-factor model. In the first one, you're trying to predict the probability that a specific policyholder will lapse. The first factor to make that prediction is a surrender charge adjusted gap between the market and the credited rate. You have a market rate minus a credited rate minus some amortization of the surrender charge. The second factor is that policyholders often react to credited rate changes independent of the market. As one of the actuaries in the audience mentioned earlier, that could be due to the equity market diverging from the fixed-income market, or it could be that policyholders are not as sophisticated and they just look at what their credited rate has done. If you drop their credited rate consistently, they're going to react against you. Policyholders don't like that. If you raise their credited rate, they'll probably react in your favor and lapse rates will go down.

You could think of the first factor, which involves the market rate, as something that agents would be constantly focused on by placing business in the market on a continuous basis. The second factor would be something that policyholders are focused on because they only see their annual statements, and they see if their rate has gone up or down. The third factor that's very important is some scaling of the policy year, which is going to be reflected mostly in the surrender charge and would be related to traditional base lapses in the older formula.

Let's say you have a model in which you have three variables and you're going to collect information on those three variables every quarter or every six months on your policies. Your goal is to fit an S-shaped curve to all the variables. In the previous formula, I showed one of them: market rate minus credited rate minus an amortization of surrender charge. For various levels of that variable, we have various lapse rates and you want to fit a function like that to your data.

Here's the actual formula that the regression will fit to:

FORMULA 1

MATHEMATICAL FORMULATION OF LOGISTIC MODEL

Probability of Lapse= Install Equation Editor and double-click here to view equation.

The formula after the exponential e should be superscript. You have a traditional regression. A linear regression would be Y equals b_0, plus b_1 times x_1, plus b_2 times x_2 and so forth. All I'm doing is transforming that. I have a probability of lapse that equals e to the b_0, plus b_1 times x_1, plus b_2 times x_2, plus b_3 times x_3 and so on for as many variables as you have over one plus E to the B_0, plus B_1 times x_1, plus b_2 times x_2 and so on. It isn't important where that formula comes from. What is important is that it will fit a nice S-shaped curve to your multivariate data. You can put your data right into a software package and out will come a lapse function. Obviously, you don't want to stop there. If you do not have data in high-interest-rate and low-interest-rate environments, you'd want to adjust this based on your own experience and judgement as to what lapses would be in those areas that are out of sample. This is a technique where you would start with your data, do a regression, and then adjust the results based on your actuarial judgement, and then monitor the new data as they become available and continue the process.

You could use the logical regression software package adjusting for your judgement and experience. Another technique is to use genetic algorithms for nonstandard formulations of the model, which is very useful.

Two very good books on this subject are D. Collett's book, *Modeling Binary Data*, Chapman & Hall, 1991, and David W. Hosmer's and Stanley Lemeshow's book, *Applied Logistic Regression*, Wiley, 1989. These books are made for practitioners. You don't need a Ph.D. to model through them. They go through how people have done studies in the past, and it walks you through the entire process of collecting data and using various software packages to get the answers.

Up until now, we haven't talked at all about path dependency or burnout, and it is very important to incorporate these into your model. This is where models get a little bit more complicated. Perhaps people have heard about burnout in the collateralized mortgage obligation (CMO) market. Basically, the definition of burnout is that policyholders who consistently ignore economic incentives to lapse should have their probability of lapsing downgraded in future periods. Let's say that people on the left-hand side of the room own annuities. Let's also say I've been giving you 3% every year for the past five years, and I will assume that none of you have lapsed. I'm going to be pretty skeptical of any model that predicts you're all going to walk out the door in the next three months. If people are always ignoring incentives to leave, you wouldn't want your model to have high probabilities of them leaving.

There are two different implementations that I'm going to talk about for incorporating burnout into your model. The first is conceptually easy to understand. You split your in-force model or your new business model into three different cells. The cells could represent hot, medium, and cold money. Then each of the hot, medium, and cold cells would have its own lapse model with different parameters. Hot money would be more interest-sensitive, cold money would be less interest-sensitive, and medium would be in the middle. You would simulate going forward with those three different models and then the make up of your business is going to change as the model develops.

The problem with that is you've just tripled the number of cells in your model. If you're constrained by run time that can be very problematic. The second way doesn't increase the number of cells.

From the Floor: If you use this method for in-force business, don't you have to use something like a pull factor?

Mr. Mungan: That's right. Yes, if you have a block of in-force business that's 10 years old, then you need some estimate as to how burned out that block is when you start the model. You can get that by analyzing any data that you might have around and, if you don't, then just give it your best estimate.

From the Floor: After 10 years the hot is gone.

Mr. Mungan: The hot money is pretty much gone. It's pretty cold.

Implementation number two is similar to the mortgage-backed securities model burnout. This can be more difficult to explain, but it is better in terms of the amount of run time in your model. I have a burnout measure, which has a ratio of two probabilities or two survival measures: a total survival which reflects all the interest-sensitive factors in your model and a base survival that would reflect none of the interest-sensitive factors. If somebody had shown that the total survival reflecting all the interest-sensitive factors was very low, the probability of still being around is only 10%. But if none of the interest-sensitive events had occurred, the probability of being around would be 70% or 80%. They'd have a very low burnout measure, and that person would be extremely burned out.

Let's talk a little bit about universal life (UL). This expands the problem to include renewal premium. If you have a UL product where you're not only concerned with a large lump of money coming in upfront, but, also, the renewal premium, how do you expand the lapse rate time modeling that we've done? This could incorporate flexible premium deferred annuities (FPDAs) sold through the tax-sheltered annuity market and regular UL products. The question is how sensitive is the renewal premium to interest rates? Basically, UL is less sensitive than single premium deferred annuities (SPDAs), as I'm sure you can all imagine, because policyholders have to be reunderwritten, and they also have to go through a high surrender charge. The general approach that insurance companies would use for UL type products is to do a thorough and complete SPDA study, and then expand that to include their UL business. Another sort of cautionary note is that the UL blocks really weren't big enough during the interest spikes of the late 1980s to make a definitive conclusion that this business is interest sensitive at all. It's less interest sensitive, but exactly how much less, I think, remains to be seen.

There are different types of predictor variables that you could use on UL for credited rate anti-selection. One factor that comes to mind is that having a strong, independent agency force is going to pressure you to keep competitive rates and a low crediting rate is going to hurt your marketing of new business. You also have to consider the responsiveness of renewal premium to your crediting rate strategy that you might have. There are, of course, different product funding levels. If you have a term type UL product, then the customer really doesn't care about the credited rate because they don't have much money in the policy anyway. From the company's point of view, they can credit pretty high rates because there really isn't a substantial amount of assets. They can get their profit component from another part of the business, such as from the mortality margin or somewhere else. Then the medium and the highly funded UL would probably be much more interest sensitive.

With persistency bonuses, there has been some work on blocks of business that had retrospective refund of charges if you persisted past year 10. This is an example of interest sensitivity or just sensitivity to basic economics that show up in these blocks of business. Perhaps part of your product gives you back 50% of your cost of insurance (COI). There have been studies where there have been large lapses of 15–20% after the bonus and very few lapses before the bonus. Policyholders were sensitive to this product feature. This was showing up on even low-face-amount ($100,000–150,000) products. It's certainly not the case that you can make a blanket statement that UL can be ignored and you just focus completely on annuities.

I wanted to make the point that you need a consistent and continuous process for monitoring your data and using the information that you get out of your customers, and then have a good annuity and UL model.

Mr. Rod L. Bubke: You talked about lapse rates, in general, and I assume you're referencing total lapsation. Have you tried applying this methodology to partial withdrawals and interest sensitivity of partial withdrawals?

Mr. Mungan: I never have personally, but that's a good point. If people have a partial withdrawal feature you'd want to find out if they've utilized it if you have any data on interest rate spikes. I haven't. I think that's something that has been overlooked.

Mr. Christian J. Shiemke: On your logistic model, do you have a minimum nonfinancial lapse rate on that or can the lapses actually go to zero, essentially?

Mr. Mungan: They can go to zero, but if you're using your data, they will bottom out at the low point in your data. If you want to make an adjustment once the model comes out of the software, you certainly can.

Mr. Jerry F. Enoch: In using the logistic modeling, what do you need to know in order to know that you have enough data to give any kind of credence to what's coming out of the formulas?

Mr. Mungan: That's a really good point. It would be totally inappropriate to just fit the model to your data without studying what kind of data you actually have.

Something that might be useful is to look at those first two variables that I presented, the difference between the market rate and the credited rate adjusted for surrender charges. You want to look at the distribution of your raw data. If you have a surrender charge adjusted gap ranging from –200 basis points to 200 basis points, with quite a few data points along the way, then you could be pretty confident that the model is going to predict that range accurately. If you have no data beyond the positive 200 basis point gap, then you want to think long and hard about what do you want the model to predict in that range, because you certainly can't rely on your data. You want to go through this process of seeing what kind of data you have in that range. How much do you rely on your formula, and then, outside of that range, you should put something in that reflects your best judgement.

Mr. Shigley: I'm going to be talking about some optionality and the stable value of options. I'm sure that most of the audience is unfamiliar with that. One of our problems is that we really don't have any good data. My first question might be, are there any good data on annuities and the extent to which policyholders react to interest rates?

Mr. Mungan: I think my answer to that is definitely, yes because of all the moving parts. If you're looking at a variable that looks at the market rate, the credited rate, and the surrender charge all lumped in together, it may be the case that your own company hasn't experienced annuity market rates of 14% or 15%. It could be following a crediting strategy that leads to large gaps between the current market rate and the crediting rate on a policy. Similarly, as surrender charges grade off a product, it creates less of a drag on lapses, and you'll see lapses pick up as the surrender charges go away. That's very similar to interest rates and a situation where surrender charges are constant and market rates go up. I think there is a fair amount of data.

Mr. Shigley: How many people in the audience are familiar with stable value funds? It looks like about 5%, so most of you are suffering from the common disadvantage, which is you really don't understand. I think this talk is designed to build some intuition that will apply to your annuity business, particularly, with respect to ripe annuities. I think you'll find that there's some intuition here that you can use.

The theme of my talk is dynamic hedging of the benefit response of option and stable value funds. A stable value fund is a special type of fixed-income option inside a 401(k) plan. Employees can move money in and out of the option at book value. In this regard, it functions just like an single premium deferred annuity. The typical stable value fund has a duration of about 2.5 years. The funds are typically managed by fixed-income managers who specialize in stable value fund management. The insurer's role is to issue guaranteed investment or synthetic wraps to these plans, which enable the plans to pay the benefits at book value.

Let's discuss the balance sheet of a stable value plan. The assets include guaranteed investment contracts supplied by insurance companies, or bank investment contracts (BICs) supplied by banks, actively managed bond funds, single security CMOs and collaterized bond obligations (CBOs), and asset-backed securities. The assets also include something usually referred to as a wrap. In this context, a wrap is an option, and it comes in two flavors: participating (par) and nonparticipating (nonpar). The nonparticipating wrap is an at-the-money option and the participating wrap is an out-of-the-money option.

Liabilities in the stable value fund are demand notes on the book value of participant account balances. These demand notes are different depending on whether the market is greater than or less than book. If the market is less than book, the demand note equals the market value of the assets, plus the put option, where the participant has the right to sell shares to the plan at book value. With the nonparticipating wrap the loss by the plan is absorbed by the GIC issuer. With the participating wrap, the loss is absorbed by the plan, unless it would cause the crediting rates to fall below zero. At that point, the issuer picks up the cost of the option.

With a participating wrap, you basically have an asset/liability mismatch. When market is less than book, the participant option is in the money and any withdrawal triggers a loss. Absent a nonparticipating wrap, the loss is absorbed by the plan. The cost of the mismatch is picked up by participants who stay in the fund.

I want to put Chart 2 into a historical perspective. This is a picture of quarterly interest rates from January 1975 through January 1996. The index rate line is the historical one-year treasury, the non-par line is the blended rate or crediting rate in the stable value fund for a three-and-a-quarter year duration fund with benefits paid on a nonpar basis. The par line is a three-and-a-quarter year duration portfolio with benefits paid on a par basis. While three-and-a-quarter is on the long side for stable value funds, it's probably on the short side for SPDA accounts.

In both cases withdrawals are assumed to react to the difference between the one-year treasury and the crediting rate. The level of withdrawals is assumed to follow a GNMA prepayment pattern, and I'll get into more detail on that later. The crediting rate or reset rate, as it's called, in both cases is based on the conventional duration reset mechanism, where the difference between book and market is amortized over the duration of the account at quarterly reset dates.

The difference between the crediting rates on the non-par line and the par line is entirely due to the par versus nonpar effect. In the nonparticipating scenario, the plan has insured the at-the-money put option and assets and liabilities are matched. In the participating scenario the remaining participants are funding the put option for the participants who leave. The losses drive down the crediting rates, which drive up with withdrawal, creating further losses until you get a classic assessment spiral.

Chart 3 shows what happens to the fund balances over the same period. The plan with participating wraps melts down because fund balances have all been withdrawn. Although it's clear that the assumptions were picked to get the meltdown, it's fair to say that this scenario is not all that improbable. The three-and-a-quarter year duration is within normal ranges. The testing period was real and the GNMA withdrawal function is not all that unreasonable based on the way some of these plans are designed.

The example was designed to raise a concern that if a stable value plan buys participating wraps, then the book value option has to be actively managed in some fashion, and that means some kind of dynamic hedging strategy is needed. The same choice exists for SPDAs. How do you manage the liability options? Do you insure or do you self-insure?

Dynamic hedging is really just a technical term for self-insurance. Table 3 is meant to show how a self-insurance program would work. The first column describes three different options, and the second column has the price for each option. The last four columns show how the price is derived. The examples are for European put options with one year to expiration on a $1,000 zero-coupon bond. The price is calculated using a Black-Scholes formula.

The first row shows that current rates are at 8%. The strike rate for the option is 10%. The price for the option is 52 cents and can be derived as a factor times the strike price, minus the factor times the bond price. The Black-Scholes formula tells you how to calculate these factors. The main reason for using Black-Scholes is it makes the factors easy to calculate and it is excellent for developing intuition. The basic idea behind dynamic hedging is that the put option can be replicated with the long position in cash, calculated as F_1 times the strike price in a short position and the security calculated as F_2 times the bond price.

The factor F_2 in the second column from the right is called a hedge ratio. It's called a hedge ratio because the replicating portfolio has the same price sensitivity as the option. By continuously changing the hedge ratio as rates move, we're dynamically matching the price sensitivity of the option. When rates move 8–10% and the option goes from an out-of-the-money option to an at-the-money option, the hedge ratio would gradually be increased from 0.04 to 0.49. The basic motivation behind the dynamic hedging is that if you adjust the hedge ratio precisely as required to reflect the current rate environment, then you should be able to replicate the same payoff as if you had bought the option or, in this case, a nonpar wrap. At least that's what happens in theory. In practice there would be trading costs and the hedging program would only work if actual volatility were equal to the assumed volatility, but for this purpose we're going to ignore that.

Let's discuss how this dynamic hedging is applied to a stable value fund or, in your case, how you could apply this to the SPDA account. The equations are shown below. Cash – HR x (MV of Assets) restates the equation for the put option from Table 3, where HR is the hedge ratio. The put option is equal to a long position in cash, minus the hedge ratio, times the security or, in this case, the market value of the fund.

DYNAMIC HEDGING OF THE STABLE PUT OPTION

1. Put Option	=	Cash - HR x (MV of Assets)
2. Demand Note Liability	=	MV of Assets + Put Option
	=	ZCB + Cash - HR x ZCB
3. Demand Note Liability	=	Cash + (1-HR) * ZCB

where ZCB is the value of a zero-coupon bond with apropriate duration

TABLE 3
DYNAMIC HEDGING: PUT OPTION (BLACK-SCHOLES)

Put Option	Price	=	F1	Strike Price	–	F2	x Bond Price
I=8%, Strike=10%	$.52	=	.0416	x $1022	–	.0420	x $1000
I=9%, Strike=10%	$4.16		.2080	x $1022		.2162	x $964
I=10% Strike=10%	$14.50		.4616	x $1022		.4922	x $929
Assumptions: Option on Zero-coupon bond; European; Time to Expiration = 1; Duration = 3; Interest rate volatility = 15%		=	Long position in cash		–	Short position in security	
			Replicating Portfolio				
$F_1 = e^{-r}(1 - N(d_2))$; $F_2 = 1 - N(d_1)$							

The second equation restates the equation for the stable value liability: the demand note equals the market value of assets plus the put option. Next, we assume that the stable value fund is invested entirely in zero-coupon bonds. In this case, it is a three-year zero-coupon bond. In the second equation, we substitute the zero-coupon bond for market value of assets. This allows us, in equation three, to restate the demand note liability as cash, plus one minus the hedge ratio times the zero-coupon bond where the cash amount and the hedge ratio are determined from an option pricing formula as described in the previous slide.

Equation three is a simple and compact expression for the stable value fund liability when market is less than book. It could also be used to describe the SPDA liability when market is less than book. As interest rates rise, the cash amount and the hedge ratio goes up and the duration of the liability goes down. If you believe that participants will exercise these options, and if you can quantify how they react, then this formula defines the required asset duration that matches the liability duration. If the price sensitivity of assets behaves exactly like this formula, then assets and liabilities are matched. Any cash outflow will fund either in a neutral position. With nonpar wraps or with purchased options, the assets will automatically match the liabilities, but if the price sensitivity of the assets is different from this, that is an asset/liability mismatch and cash outflows will leave the fund in the loss position. To avoid losses with participating wraps, the asset portfolio should, therefore, be managed to the same duration as this liability. Equation three is a key insight behind self-insuring the put option.

I'm going to spend the rest of the time showing how this gets applied in practice. To apply this to real stable value funds or SPDA accounts, you need to first take a closer look at how these options would get exercised. In practice, even though participants own at the money options, they tend not to exercise these options until they're way in the money. Thus, if the blended rate or crediting rate were 7%, most participants might wait until current rates were up to 10 % before they were sufficiently motivated to exercise their options.

Chart 4 illustrates the difference between a conventional at-the-money option and these participant options. The payoff for the conventional at-the-money options starts at 7%, or point Y, and increases as the rates rise. In contrast, a stable value payoff profile stays at zero until rates hit a much higher threshold. In this case, it's assumed to be 10%. But after rates hit that threshold, the payoff jumps up at the vertical solid line until it's equal to the at-the-money option with a strike at 7%.

Technically, this payoff profile is the sum of two simpler components. The first component is a digital option that pays a flat amount equal to the vertical solid line whenever rates are greater than the 10% strike. The second component is a conventional out-of-the-money option with a strike at 10%.

When we put this modified payoff pattern through a Black-Scholes pricing model, we get a result as shown in Table 4. The exhibit shows price, hedge ratio, and net portfolio duration for the modified option in Chart 4. Let's focus on the first row. When current rates are at 7% the option, which is not exercised until rates hit 10%, costs only 16 cents. That's less than two basis points. The hedge ratio is 0.02. If we wanted to replicate this option inside a stable value plan, we would need to shorten duration from three to 2.94 years to match the price sensitivity of the three-year duration plan with a nonpar wrap or purchased options. This is not very dramatic.

TABLE 4
DYNAMIC HEDGING:
MODIFIED PARTICIPANT PUT OPTION

PUT OPTION	PRICE $	PRICE %	HR	PORTFOLIO DURATION CASH + (1-HR)* ZCB
I = 7%	$.16	.02	.02	2.94
I = 8%	3.97	.41	.28	2.16
I = 9%	21.55	2.32	.91	.27

Assumptions: Option on Zero Coupon Bonds
European; Time to Expiration = 1;
Duration = 3
Interest Rate Volatility = 15%
$F_1 = e^{-r}(1-N(d_2))$; $F_2 = 1-N(d_1)$

The third row shows the blended rate or crediting rate. We're still only at 7% and the current rate will now be up to 9%. The option, which has the modified payoff that pays the same amount as an in-the-money option at 7%, but has a delayed exercise until the strike is at 10%, is now worth $21.55 or 2.32% of the fund, and the hedge ratio is 0.91. If this option were properly hedged or self-insured inside the plan, duration would have to be moved from three all the way down to 0.27 to match the price sensitivity of a three-year duration plan with a nonpar or purchased option.

These examples are still a little artificial, because they all assume that 100% of participants would all react at the 3% rate differential and none would move before that. In the real world, different participants are likely to react at different benchmark levels, and thus we're likely to see participants move at progressively greater speeds as rates move progressively higher relative to the blended rate or crediting rate. At a low rate like 5% per year, when the blended rate differential was 2%, maybe 10% will transfer when the crediting rate differential is 3% and so on.

In order to actually manage the benefit risk we need to find a behavior function. The big question then is do we have any guidance on what this behavior function would look like? As a working hypothesis the GNMA prepayment function is a good starting point for addressing this question. GNMA prepayment speeds are based on real experience with option efficiency derived from real people who are largely the same people who invest in stable value plans or that purchase SPDAs.

The right-hand side of Chart 5 has a picture of a regression of GNMA prepayments expressed as a function of the spread between the refinancing rates and mortgage coupons. On the left is a table that roughly translates the regression into a stable value context. The table relates crediting rate differentials to withdrawal rates. For example, when the crediting rates are 2% lower than current rates, then withdrawals from the stable fund will flow out at an annualized rate of 16.5%.

Withdrawals will be lower for smaller differentials, and withdrawals will be higher for bigger differentials.

We don't really have any direct evidence for the level or magnitude of this behavior function for stable value plans. The last time interest rates were significantly higher than blended rates was in the early 1980s when most participants didn't have many other choices for where to put their money. We really don't know whether the behavior function in stable value plans will be at 100% of the GNMA level, or 50% of the GNMA levels, or at some point above or below.

We'll focus on the summaries at the bottom of Table 5. It develops option prices and hedge ratios, assuming that participants will react at 50% of the GNMA function or at half the levels in Chart 5. I think that's not an unreasonable assumption for some of the plans that we're dealing with.

The left-hand side of Table 5 sorts participants according to their respective exercise strike threshold, and the percentage that's assumed to strike at each point. If we focus on the fourth row of the first block, we see that 14.2% of the participants will move their balances if rates go to 10%. The crediting rate throughout the exhibit is assumed to remain constant at 7%. The 14.2% that moves at 10% would be reacting to a 3% differential between current rates and the crediting rate at 7%.

The other blocks have prices and hedge ratios for various put options at each strike rate from 7% to 12%. The entries in the second block assume current rates and crediting rates are both at 7%. In the third block the crediting rate is at 7%, but current rates are at 8%. In the last block, the crediting rate is still 7%, but the current rate has moved to 9%.

TABLE 5
DYNAMIC HEDGING IN PRACTICE PRICE AND HEDGE RATIOS-50% GNMA

		Current Rate at 7%		Current Rate at 8%		Current Rate at 9%	
Weight	**Exercise Begins**	**Price/$1000**	**HR**	**Price/$1000**	**HR**	**Price/$1000**	**HR**
1.00%	7.00%	$11.38	.49	$31.77	.80	$58.08	.94
2.81%	8.00%	6.93	.40	27.00	.84	55.29	1.01
8.23%	9.00%	1.65	.14	14.27	.67	42.90	1.13
14.20%	10.00%	.15	.02	4.15	.27	23.21	.92
18.21%	11.00%	.01	.00	.62	.06	7.98	.45
18.21%	12.00%	.00	.00	.05	.01	1.70	.13
Weighted Total		$.47	0	$2.98	0.1	$10.72	0.4
Dynamic Stable Value Duration			2.91		2.7		1.7

Prices and hedge ratios in all these blocks are for the modified payoff structure where the option is in the money at 7%, but where the exercise is delayed until rates rise to the strike point. At the bottom of each block is the weighted average option price and the weighted average hedge ratio for the package of options in each scenario, and this is what we want to focus on. In the second block, where the crediting rate and the benchmark rate are both at 7%, the cost of the weighted average option package is 47 cents and that's only five basis points. That's excluding expenses and cost of capital. The hedge ratio is 0.03.

One way to interpret this would be to say that when book equals market, a plan which self-insures the put option would have to move duration from 3 to 2.9 to match the price sensitivity of a three-year duration fund with purchased options. The situation changes when current rates move to 8% and the blended rate or crediting rate is still at 7%. The weighted package of a one-year option cost is now 30 basis points. The hedge ratio goes to 0.14. A plan that self-insures would move duration from 3 to 2.58 to match the price sensitivity of a three-year duration fund with purchased options. Finally, when the current rates are 2% higher than the crediting rate, then under these assumptions, annual option costs and wrap fees would rise to 107 basis points, and the hedge ratio increases to 0.37. This means that duration for the self-insured plan would have to drop from 3 to 1.89 years to maintain the same price sensitivity as a three-year duration plan with purchased options.

The exhibit does a nice job of building intuition about the cost of participant disintermediation. It shows that when market is close to book, the benefit responsive option is very cheap. Under these assumptions when book equals market, the option charge is about five basis points. But if rates rise, the price of the benefit responsive option goes up geometrically. Furthermore, the increase in costs as rates rise is much steeper than the decrease in costs if rates decline. To the extent rates are equally likely to go up as they are to go down, the expected cost for the ensuing period in the statistical sense will be higher than the five basis points calculated in the initial period.

I'm going to end at this point and hope that these remarks have given you some ways of applying the cost of the options in the annuity market. I think it has applications to durations that are selected for SPDA accounts, and I think the applications are primarily for situations where your annuities are ripe, and there are no longer significant surrender values.

Mr. Mungan: When you talked about managing the duration in that way, how is that done in practice? Do you use futures contracts or are you actually trading bonds?

Mr. Shigley: You have two choices. You can buy the options or you can manage duration. In this particular case, the people who are managing the duration are the stable value managers, and they can choose to buy the option from us or to manage it themselves. To my knowledge, they're really self-insuring, and they're not managing the duration at all.

Mr. Scott L. Fitzpatrick: The title of this session was "Bad Behavior: Policyholders and Interest rates," so I think I can direct this to you, Ken. You showed us how policyholders behave, but I don't see how that's bad behavior. If there's no value in what the product's delivering, they should pull their money out. The only reason they wouldn't be getting value would be because either the expenses are too high or the investment strategy is not returning what it should, right?

Mr. Mungan: That's right. As somebody who was asked to be on this at the last minute, I can say I didn't name the session. Obviously, policyholders are going to act in their best interests and they should. Insurance companies are acting to give their policyholders advice and access to markets, and if they're not meeting their policyholders' needs, the policyholders are going to walk out. My objective has always been to have just a clear estimate of what the policyholders are going to do, and then to come up with the best investment strategies given that behavior.

From the Floor: Whose behavior is being modeled? We're talking about participants and they pretty much act on their own behalf. With policies, are you trying to regress the advisor's behavior or the policyholder's behavior and to what extent is the data reusable? I'd like a comment on that.

Mr. Mungan: That's a really good question and that's why I always stress that you need to do this process for each distribution channel and product separately. Because if you have a distribution channel, say a broker channel where the advisors to the policyholders are very active in studying the market and they know that they can generate a commission if the company is not competitive in its product and its crediting rate strategy, they're going to be advising the policyholders to move their money to something that's more appropriate for them.

In the model that I showed that had the factor of a gap between the market rate and the credited rate, you could think of that as something that is really going to be driven by an agency force that's actively placing new business in the market. But if it's something where the policyholders are pretty much on their own, then they may not be so aware of the market. They're going to react more to how their credited rate goes up or down. They're going to react to the *U.S.A. Today* effect; if there's a big story about the security of insurance companies in the paper, or on the news, or if Mike Wallace shows up with a microphone pointed at Klaus, that's going to affect policyholders. You want to come up with any predictive factors you can that will segregate the different behaviors between policyholders and agents.

Mr. Shigley: There was a comment earlier about using stock market returns as a variable to model the policyholder behavior. We actually took a look at some stable value plans. We did regressions on withdrawals from plans against stock market behavior. We found that there were a few participants who reacted very heavily to changes in the equity markets. We found that the R-squared for stock returns on participant behavior was about 0.14. The stock market had a very strong influence on a very few participants. We found that it didn't help explain a lot of activity, it explained a little of the activity very well. In contrast to that, the behavior in GNMAs has a R-squared of around 0.89. There's a much stronger relationship between interest rates and mortgage refinancing than there is between stock market returns and participant disintermediation.

Mr. Ziegler: Ken, your formula for the burn out reflected the difference between total lapse rate and base lapse assumption. I think that has some good points to it. I think it misses the point that you made a little earlier about the agent involvement. Typically, when policies get fairly old, there's a higher proportion of orphans. If you've lived through a pretty good lapse period, the difference between total lapses and your base lapses is pretty small, that ratio is about one, Those policyholders, because of the lack of agent influence, probably are less sensitive. It is probably worthwhile to make some reflection of that in a burnout factor.

I wanted to make a comment, also, on Dan's point about the stock market. We've tried to look at our single premium deferred annuity lapses. When we're trying to validate and come up with factors for the formulas for excess lapses, we've tried to strip out lapses somewhat attributable to the stock market which is certainly a less-than-perfect science of doing that. There is a key piece of information that has helped us in that; again, it's not perfect, but it might be something that other companies are doing or want to comment on. We have a conservation group effort that calls lapsing policyholders and we generally find out where the money's going. We can, over time, look at what percentage of the money is going to either companies or products that we think are generally equity based as opposed to interest-rate based. That percentage does vary over time, so we've tried, in a crude way, to strip those lapses out before trying to fit the interest rate thing.

Ken, you provided a formula that has the ability to build in a lot of different factors. Have you tried to build in an equity component to that as well, because I think that would help solve that kind of a problem?

Mr. Mungan: That's a good point. That whole process I described, called the lapse-analysis process, is something where you can always be trying out new variables. To address your first point about orphan policyholders, I'd say you could have a binary variable. It is 1 if the policyholder still has an agent, and 0 if they have no agent. That is something that I have seen that has a definite effect and that goes right along with account value. The higher account value SPDAs are going to get more attention by agents. If there are lower account value policies with no agent currently assigned, that money tends to be pretty cold. To manage your in-force business, if you're trying to determine what the impact of the change in crediting rate strategy is going to be on a block of in-force business, you should have a lapse formula that has as many factors as you can come up with that your regression tells you are significant in terms of determining policyholder lapse rates.

That sounds like a decent approach to taking the stock market into account. You can strip it out and then have some way of putting it back in when you model scenarios in the future in terms of the average amount going into the stock market or you can just leave it in. But you should definitely get a feel for how the stock market is impacting your policyholders. It could be, as Klaus said, that it has very little impact on a lot of people, because they simply are unwilling to accept the risk of the stock market. I know a lot of people will stick with nice, stable, insurance-company-fixed products rather than take the risk of losing their principal. On the other hand, some people in the hot money category will jump over immediately. By having another variable in your lapse formula, either if it's a logistic formula or a more traditional one, you can capture that.

Mr. Shigley: What is the average duration for fixed-income SPDAs? As a contrast, in the stable value industry, the average duration of a stable value plan is about 2.2 years. My guess is that the fixed-income SPDAs tend to be a duration of probably at least a four. I'm not sure exactly what that duration is, but I see the two products as being very similar. The risk profile in the two is enormously different. I have these theories that say what is being managed or what is being optimized are different things. I'm not sure exactly what's being optimized for SPDAs. It tends to be managed more aggressively than stable value funds. I'm not sure what, if any, facts there are. Do you have any facts?

Mr. Mungan: Before I go on with my duration answer to Klaus's question, I'd like to know whether anyone in the audience actually calculates duration on SPDAs? Could you tell us about it and how you carry out the technique?

From the Floor: We discount the present value and check the interest rates a little bit and we calculate on that.

Mr. Mungan: Are you discounting pathwise or discounting at a constant rate?

From the Floor: Pathwise, yes.

Mr. Mungan: That's a good technique. So you have many scenarios and you're discounting each scenario pathwise and then shifting up and down. I've seen durations in the four to five range. Is that consistent with your durations?

From the Floor: Ours are usually four or a little bit less.

Mr. Mungan: If you have a long surrender charge product or high surrender charges, it is going to have a longer duration.

Mr. David M. Walczak: We have a block of old SPDAs without surrender charges that are basically current interest or new money spreads. We also have a big block of two-tiered annuities with permanent surrender charges. Our option adjusted duration or, as you would call it, pathwise shift discounting measure of effective duration, is under two for the surrender charge gone, fast credit rate change block, and somewhere in the four to six range or higher depending on how you model the options for the two-tier permanent surrender charge block.

Mr. Mungan: Sure.

Mr. Lone-Young Yee: We calculate duration for our deferred annuity block at my company. The key component of the duration is the interest rate credited strategy. If you have a short-term interest crediting strategy, your duration will be very, very small. You're going to state a certain percentage no matter what. Duration could be about four or five years. If you're going to reset interest rates at 50% of market changes, your duration will come down probably back to two.

Mr. Mungan: That's an excellent point. Policyholder lapses are part of a dynamic process and the interaction between investment strategy, lapse rates, and crediting strategy. All of those factors interact when you're calculating duration. If you're investing in a strategy that allows you to have crediting rates after the surrender charge, which respond well to the market rate, so you're not stuck with low crediting rates when the market rate spikes up, then your duration can extend out. If you're following a deliberate crediting strategy of ratcheting crediting rates down, as the products do during a surrender charge, then given the policyholder reaction factor that I talked about (where if you lower the crediting rate, policyholders are going to see this and they're going to be out the door), that's going to lower your duration on the product. As you're developing your asset/liability modeling strategy, setting up the best investment strategy and crediting strategy, you need to consider the entire picture. Klaus brought up a point in terms of what you are optimizing when you're developing that strategy? One thing that I've worked with in the past is some kind of risk-adjusted mean present value measure where you recognize, as an SPDA writer, that you're going to retain risk no matter what. It's not in your best interest, as a company, to lay off all of the risk. Of the risk that you retain, you look at the shape of the distribution and then have some measure of the downside risk and try and balance the mean return that you're going to get with the amount of downside risk that you're maintaining.

CHART 1
PERCENTILE GRAPH FOR LAPSE RATES—SPDA BLOCK

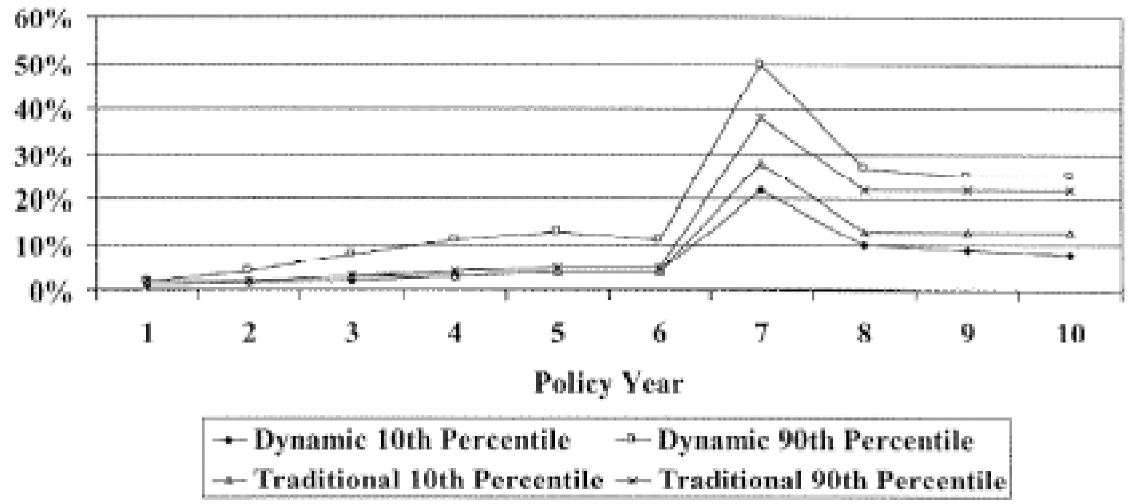

CHART 2
CREDITING RATE COMPARISON: PAR VS. NON-PAR
(CONSTANT DURATION=3.25; BEHAVIOR=100% GNMA)

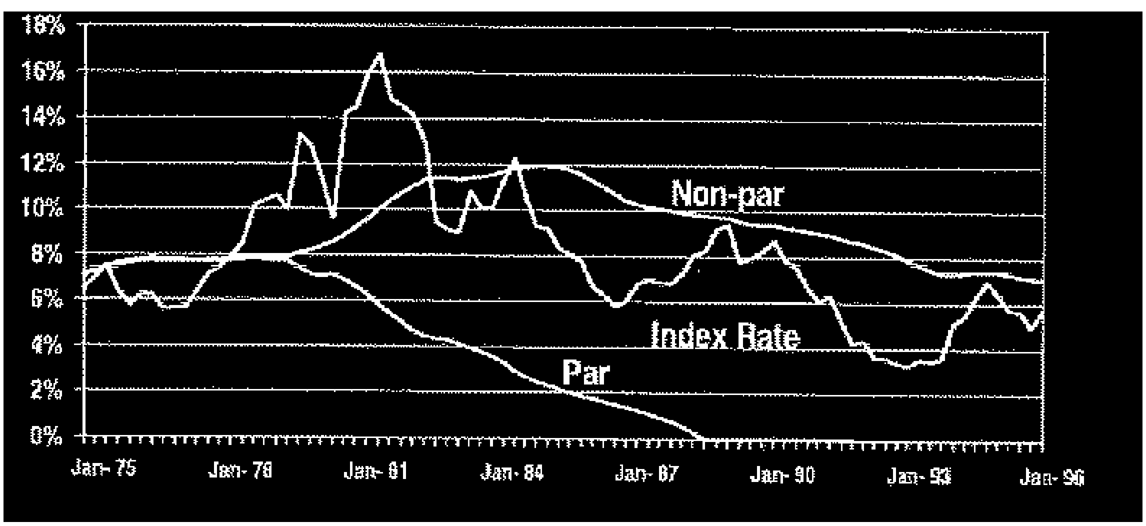

CHART 3
FUND ACCUMULATION COMPARISON: PAR VS. NON-PAR
(CONSTANT DURATION=3.25; BEHAVIOR=100% GNMA)

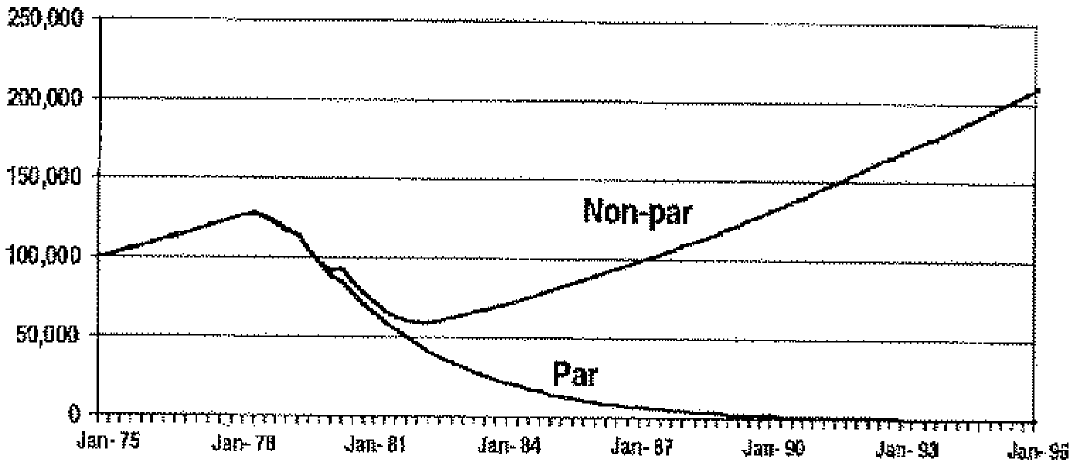

CHART 4
PARTICIPANT PUT OPTION PAYOFF PROFILE

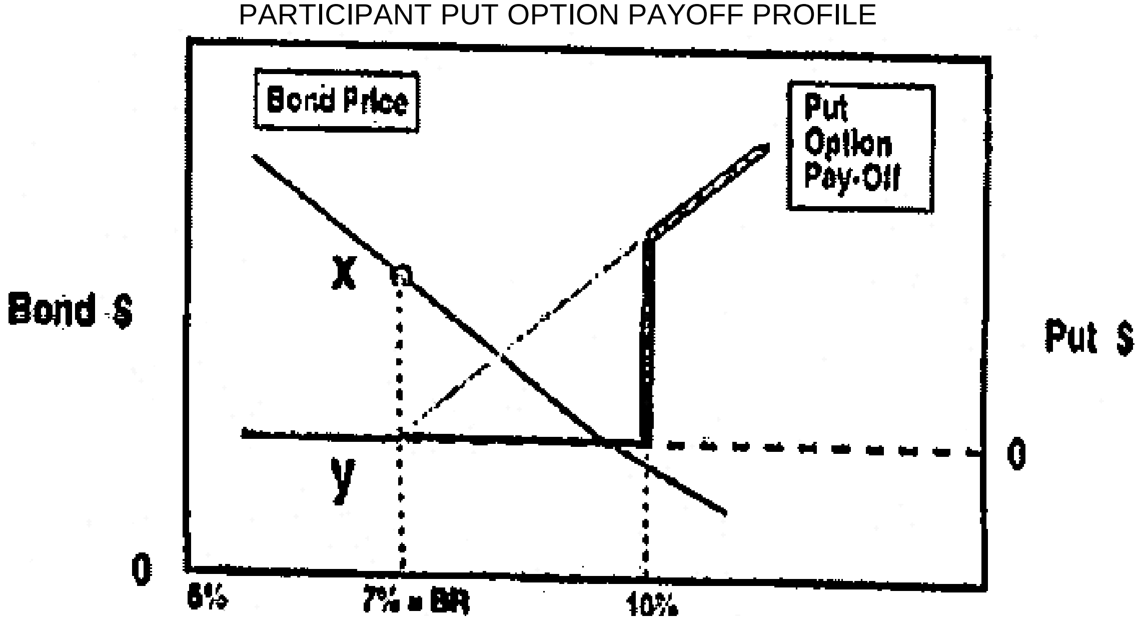

CHART 5
DYNAMIC HEDGING—BEHAVIOR FUNCTION
GNMA PREPAYMENT RATES
COUPONS 6.0–11.0
AUGUST 1990-NOVEMBER 1995