RECORD OF SOCIETY OF ACTUARIES
1989 VOL. 15 NO. 2

ANALYZING THE ASSETS FOR ASSET/ LIABILITY MANAGEMENT PURPOSES

Moderator: JOSEPH J. BUFF
Panelists: SHELDON EPSTEIN
GREGORY D. JACOBS
Recorder: JOSEPH J. BUFF

- o This session will discuss the various asset related items that are needed to perform necessary projections.
  - -- Asset in-force data
  - -- Interest rate scenarios with associated probabilities
  - -- Asset default and/or call or prepayment assumptions
  - -- Quality spreads
  - -- Liquidity
  - -- Investment strategy for future cash flows investments
    - o Spreads
    - o Asset features
    - o Maturities
    - o Liquidations

MR. JOSEPH J. BUFF: The title of my relatively nontechnical talk is "Megabytes and Methodologies." The theme is the advance of inexpensive computer power and the concomitant advance of professional practices and actuarial methodologies. A better understanding of this interdependence ought to lead to better professional success and satisfaction for all of us. The state of the art of computers and the state of the art of actuarial science are not completely separate. My thesis is that actuaries are doing things they didn't used to do simply because more powerful and user-friendly hardware and software are available now. This trend has been going on for years and it will probably continue to go on for many years to come. I am going to take you through a personal retrospective on computing power and actuarial science.

Ruskin said that "the work of science is to substitute facts for appearances and demonstrations for impressions." A lot of the technical work that many actuaries do is oriented to facts for appearances and demonstrations for impressions. We calculate and project financial variables to render more concrete and specific the vague and intuitive issues surrounding the financial security business. Peter Drucker said somewhere in his extensive writings on management science that "if you can't measure it, you can't manage it." This is a practical counterpoint to Ruskin's cogent observation. I think they are saying the same thing. I would go one step further and make a quotable statement myself: "If you can't compute it, you can't compute it." This truism may be worthy of the great thinker Yogi Berra, who once said "the ballgame isn't over till it's over." But Yogi has a point and I have a point. It is not possible to make use of technical methods and formulas unless it is feasible to execute the calculations called for by those methods and formulas.

Actuaries play an interesting role in the economy because they form a bridge between science and industry. Actuaries practice actuarial science, and at the same time they help run businesses. Insurance and pension managers, regulators, consultants, even academicians, are business people. As business people, we are dependent on computer hardware and software for data processing, for financial analysis, even for getting out our own paychecks on payday. Hardware and software facilities are completely dumb (i.e., unable to speak) without professionals determining and programming the methods and formulas that the equipment is supposed to run. For actuaries, many of those methods and formulas deal directly or indirectly with pricing and reserving. This is equally true for pension actuaries, property/casualty actuaries, or life actuaries.

Actuaries, therefore, are dependent on computer technology. Could computer automation reduce the need for actuaries? Could computers replace actuaries? Well, let's take a look at two trends over the last twenty years. Trend one is the capabilities and costs of personal interactive computer power. Trend two is the complexity and challenge of the job actuaries fill.

In 1970, a "pocket calculator" was a slide rule. For those of you who have forgotten what a slide rule was, or in case you are too young to have ever seen one, they have a number of different scales on them. You slide the middle thing back and forth. Then you moved the little window with the crosshair, called a cursor, to the right spot and read off the answer by eye. With practice you could get answers to three significant digits. You didn't have to worry about the batteries running out. This thing was powered by hand. However, there were potential problems with slide rules. As mechanical instruments, they could fall out of adjustment. If the scales became misaligned you would get the wrong answer. There was a debate going at the time whether the best upmarket slide rules were made of steel or bamboo. Steel you see, was immune to moisture but relatively sensitive to temperature distortion. Bamboo, on the other hand, was more stable when temperatures changed, but bamboo could absorb moisture and loose accuracy under humid conditions. In fact, on a rainy day a bamboo slide rule could jam up. I guess before people were able to give the excuse "the computer is down" when something got delayed, they could say "my slide rule is stuck!" Good slide rules cost about $30 in 1970. By the way, the cost figures here are not inflation adjusted; 1970 dollars were worth almost three times 1989 dollars.

In the early 1970s, electronic calculators started to become a consumer item. The original pocket-sized calculators were barely pocket sized, since they were pretty thick and heavy. They were quite an advance over the slide rule. In 1973, still in college, I went with my girlfriend of that time to buy her younger brother a present. The present was a four-function calculator. It did addition, subtraction, multiplication, and division. It had an eight-digit output display. It had no storage memory. It was on sale and it cost $150. When we gave it to the young man, he became the first student in his entire school to own a pocket electronic calculator.

This early generation pocket calculator was thick and heavy in part because of its power source: a pair of AA batteries. The calculator used LEDs, light-emitting diodes, for the output display. These glowed in the dark. They took a lot of power, hence the heavy batteries. In continuous use the batteries lasted about four hours, I think. Then you needed new batteries.

Time marches on. It is now 1977 and I was in graduate school. Now we would get "scientific calculators." These hand-held devices could compute all sorts of logarithmic and trigonometric functions. Some of them were even programmable, and they had storage memories. They cost about $150. So pocket calculator prices had held steady in nominal dollars, while their capabilities had increased substantially.

Better output display technology was on the drawing boards, but I don't think it was in the stores yet. Liquid crystal displays, LCDs, were on the way. These relied on ambient lighting so they used very little power. Pocket calculator power sources would transition to hearing aid batteries that lasted for more than a year. This led to a breakthrough in size and weight. Credit card-sized calculators would eventually appear. Recently I bought a couple for $4.95 each.

But a social problem had cropped up in the late 1970s. Not every student could afford a good pocket calculator. To many people, especially those in college, $150 was a lot of money. Universities were becoming concerned that the emergence of pocket calculators gave an unfair advantage to the rich. They could do their homework faster, and they could do better on exams where slide rules and calculators were permitted. I'm not sure this problem ever was solved.

Now it's 1980 and I'm an actuarial student. This was the dawn of minicomputers. These were tabletop computers. We had one in the group department of my employer at that time. It had 24,000 bytes of core. Remember that figure 24,000 as we go on. The thing printed output on a dot matrix printer which was slow and very noisy. It also had a screen monitor, which had pretty low resolution by today's standards. This machine could crank through the quarterly accruals of dividends on all the group cases in about twelve hours. We used to set up in the afternoon and collect the results. We were OK as long as the paper feed didn't jam. These things cost $21,000 new. Starting salaries for actuarial students with one exam in New York City were less than $20,000 at the time.

In 1985, we were well into the modern era of mainframe timeshare computing. This provided a truly interactive environment for heavy duty calculations and big file manipulations. The system

I used at the time gave you 100,000 bytes of core, quite an improvement from that minicomputer in 1980. Now you could do very complex financial calculations in a rapid time frame. The one drawback of this computing environment was cost. The mainframes were expensive. There typically were usage chargebacks every time you signed on and did work. Sometimes these chargebacks were paper allocations only, if your employer owned the equipment. But the costs did show up on your department budget and we were always concerned to keep the costs down.

Back to the present. In 1989 we have personal computers (PCs) that are pretty powerful. A standard device now has a core of 640K bytes, 25 times what you got in 1980. Also, current machines come with individual hard drives that offer random access storage of 70 megabytes, even more. Modern PCs use laser printers that are fast, quiet, and produce letter-quality printing. Speaking of asset/liability management, you can now run extensive calculations for scenario modeling or duration analysis in just a few minutes, right in your office. Depending on how portable you want the PC to be, prices fully equipped for the hardware and operating system run about $5,000-10,000. There is a lot of good software available.

So what's been happening? Where have we been going for the last twenty years? The partnership between men and machine has been thriving. We have seen major advances in actuarial science, and major advances in computer hardware and software. Without computer advances, we would not have had actuarial advances. "If you can't compute it, you can't compute it."

With all these gradual changes, what does it all mean to us as individual professionals? Are our jobs getting simpler? Are we in danger of being replaced by computers? Well, let's take a look at a couple of aspects of actuarial responsibilities and how they've changed in the last decade.

Let's take a look at the technical calculations called for by state insurance departments. We've always had to comply with regulations about pricing and reserving. In 1979, when I was an actuarial student, one of my jobs was to do the periodic reports of individual health loss ratio experience for the State of New Jersey. The first couple of times I did it, it seemed a complicated process. In 1989, one of the things many of us have to do is scenario testing for New York Regulation 126. This is a vastly more complicated process. The loss ratio report was done by hand on a spreadsheet, using an adding machine. Regulation 126 needs to be done on a mainframe or current-generation PC. It is a challenging responsibility to complete a Regulation 126 filing and sign off on it as the qualified actuary. If we didn't have desktop computers with 64OK cores and affordable cash flow software, we wouldn't have a Regulation 126. Presumably then, we also couldn't use the higher valuation interest rates in New York State.

Let's take a look at another job role. Namely, investment strategies and making investment assumptions for actuarial calculations. In 1979, everybody did asset shares, not scenarios or durations. For an asset share, we used a simple assumption about after-tax interest rates. That was the extent of our asset/liability management. Now, we're here in Vancouver in 1989 learning about option-adjusted durations of assets and liabilities. That's quite a change. If we didn't have more powerful computers than ten years ago, we probably wouldn't have dared think about calculating the durations of our company balance sheets. Now some companies take this technology for granted in helping run their companies.

So actuarial roles have evolved as computers have evolved. Our jobs are more complex, more challenging, more rewarding than ever. In the 1950s, when early mainframe computers started to appear in the workplace, many pundits predicted that computers would replace middle management in major corporations. You can find it said in books and articles in the 1950s that the role middle managers played in allocating resources and making tactical decisions would be done better by computers. Well, it hasn't happened yet. I doubt it is going to. Have modern computers reduced the need for actuaries by automating the financial security workplace? Again, it doesn't seem to be happening. Demand for actuaries is as strong as ever.

Let's look ahead to the 1990s. What can we expect as actuaries, given the dependence of professional practices and methodologies on the state-of-the-art of hardware and software? Computers will continue to get faster and cheaper. Software will become more powerful, more versatile, easier to use, and less expensive in real dollars. These trends will let actuaries become even more aggressive in the technical methodologies they are able to develop and apply in their work.

In general, we will see management information systems becoming more powerful and more effective. We will do a better job of substituting facts for appearances and demonstrations for impressions. We will manage better because we will measure better. It seems quite likely that in the year 2001 actuarial roles will be as challenging and as satisfying as they were said to be in a recent U.S. Department of Labor survey, namely the best job around.

MR. GREGORY D. JACOBS: The topic that we are going to be talking about is Asset/Liability Analysis and Projections, specifically from the asset side. There are three key areas that I am going to be talking about, again, from the perspective of asset/liability projections (the asset side). I am first going to talk about the existing asset detail that you need, and I am going to go into a great deal of detail for collateralized mortgage obligations (CMO), because those seem to be the most exotic assets that are out there. We won't spend a lot of time on callable bonds or straight bonds or anything like that. The next area that I am going to get into is the two or three critical asset assumptions that we need to deal with from an asset perspective when we tackle an asset/liability analysis. And finally, I will talk briefly on asset strategies.

When you tackle this, the first thing you have to do is figure out what types of securities you have (presumably you are dealing with a book of existing assets). Here is a laundry list of the key fixed-income sort of assets that are probably in your portfolio and the key items that you are going to need to get out of them.

For bonds, you will have to know if it is a bullet bond, a sinking fund, or a serial bond. You will have to know if it's callable and whether or not it's convertible. For mortgages, you will have to know if it's fixed rate or variable rate, the amortization period, and the structure. A lot of commercial mortgages are 15-year or 30-year amortization with a 5-year bullet or a 5-year balloon payment. You need to know how that amortization structure works. It's not intuitively obvious if you look at just the information that is contained in your annual statement. You need to know the prepayment provisions, one of the critical asset assumptions that I will talk more about. And finally, the newest category of assets that seems to be big in insurance companies' books -- CMOs. I will get into a fairly detailed example of how we do the modeling for those.

Asset-backed securities include car loans, credit card loans, and some of the stuff that banks have that they would like to securitize and distribute. A lot of companies have been investing in those. You need to know the same sort of information: the underlying asset, what is it that you are essentially buying into, and the other information that I have listed here: fixed versus variable rate, amortization periods, prepayment provisions, etc. The other kind of fixed-income security that is generally on people's books is private placement. Each individual private placement is unique. You need to get into the files of the investment department's records to find this information, again, because it is not intuitively obvious from the annual statement. Finally, there is common stock, preferred stock, real estate, options, futures.

In most of the situations that I work on, most of the lines of business that we are analyzing are supported by fixed income securities, so that those first four items, together with cash and any treasuries that you have, are the ones that you need to deal with.

A CMO starts out with an underlying asset of a mortgage. We'll use a simple example (Exhibit 1).

EXHIBIT 1
CMO EXAMPLE

Payment	Interest Portion	Principal Portion	Total Payment	Yield Curve
1	$100.00	$163.80	$263.80	7.51%
2	83.62	180.18	263.80	8.51
3	65.60	198.20	263.80	9.51
4	45.79	218.02	263.80	10.51
5	23.98	239.82	263.80	11.51

Price $1,000; Yield 10%; Average Life 3 Years

This is a 5-year mortgage for $1,000 at 10% with annual payments at the end of the year. This is the normal principal and interest structure. The annual payment -- interest and principal -- is

$263.80. The price of this bond is $1,000. The weighted average life based on cash flows is 3 years and the yield is 10%. Now what a CMO does is, it breaks it into tranches. This example, hopefully, will show how it works (Exhibit 2).

Tranche 1 is the first 2 principal payments, where we will be getting the first two principal payments and the interest that is attached to them. That's what is shown in column 1, the items under tranche 1. In tranche 2, I don't get any principal payments during the first two payments. I only get interest. That is why I am getting the $41.62. After tranche 1 has been fully paid off, then I start getting principal payments until my tranche 2 is totally paid off. Finally, I end up with tranche 3 which is my final principal payment. During the first four years I am only getting interest. When tranche 2 is all paid off, then I complete the deal and I end up getting my principal payment on my tranche 3.

The sum of these three tranches is exactly the mortgage shown in Exhibit 1. Now, in the marketplace, based on that yield curve shown earlier, the price and the yield for each of these is shown at the bottom. We can take a 10% earning asset with an average life of 3 years and split it in the marketplace for a 1.5-year instrument at 8.16% or a 3.2-year instrument at a little under 10%, or a longer term 4.3 at 11.11%, just by manipulating how the mortgage principal payment structure works.

In graphic form, Exhibit 3 is what a CMO looks like. The mortgage payments are level. Again, I am assuming this is a fixed rate and I haven't factored in any prepayments yet, but the mortgage payments are level. In the first tranche there is a little bit of interest above the curve line as the interest portion of the payment, and below the curve line is the principal payment of the mortgage. In the first tranche you get a little bit of interest and a lot of principal; in the second tranche you get interest in the first couple of years and then you start paying off the principal; and in the third tranche you get interest for four years, then you start paying off the principal.

But leave it to the Wall Street community to complicate matters by creating a Z-bond (Exhibit 4). The Z-bond takes the interest payments that were going to the old tranche 3 ($23.98), and we plow that back in to prepay the principal in tranche 1. So, it is going to make tranche 1 and tranche 2 be prepaid quicker by taking the interest portion that would have been assigned to tranche 3 and moving it into principal payments in the earlier tranche. All of the other numbers fall out of that prepayment structure coming through the tranche.

The end result of this is: We have a shorter and a lower-yielding tranche 1, and we have a longer and higher yielding Z-tranche. Normally, these things are 15- or 30-year instruments, so the Z-tranche generally has an average duration of about 25 years. This example is a little exaggerated. You don't think of a 4.7-year instrument as being long term, but in a normal CMO arrangement a Z-tranche probably has a 22-25 to 27 year time frame, reasonably high yielding, and very volatile because there is no cash flow coming in the early years. It is almost like a zero coupon bond, hence the Z in Z-bond or Z-tranche. So in the marketplace, you take a 5-year mortgage and it gets split up by the Wall Street types into these three securities. That is what you as insurance company investment people or actuaries analyzing the investment portfolio are looking at. If you are buying tranche 1, they are reasonably easy to predict. If you are buying tranche 2, you have to know what is going on in tranche 1 before you can predict tranche 2, and so on. I'll explain that in more detail as we get into prepayments (Exhibit 5).

What's going on in this example is that I assumed there is $100 a year prepaid; that instead of the regular $260, they are going to pay $360. What goes on with prepayments is, under the principal column of tranche 1 you will see an extra $100 goes to pay off the principal. This has a tendency to shorten each of the average life times of the tranche, and all of the remaining principal balances on tranche 1 are affected. Then tranche 2 starts paying a little bit earlier. You can see now that the Z-tranche is only 4 years. You have to wait only 4 years to start getting your principal and interest payments instead of what was originally a 5-year mortgage.

Exhibit 6 will show you the importance of looking at these things in your portfolio. Before prepayment you could have either bought tranche 1, 2 or Z, and with these parameters associated with it -- the average life and the yield rate. With $100 of prepayment you can see some interesting things happen. One, all of the average lives are shortened. The longer the original life, the more dramatic the shortening is.

CMO MECHANICS

	Tranche 1		Tranche 2		Tranche 3	
Payment	Interest	Principal	Interest	Principal	Interest	Principal
1	$34.40	$163.80	$41.62	$	$23.98	$
2	18.02	180.18	41.62	--	23.98	--
3			41.62	198.20	23.98	--
4			21.80	218.02	23.98	--
5					23.98	239.82
Price	$352.69		$417.47		$229.82	
Yield	8.16%		9.89%		11.11%	
Average Life	1.5 years		3.2 years		4.3 years	

EXHIBIT 3

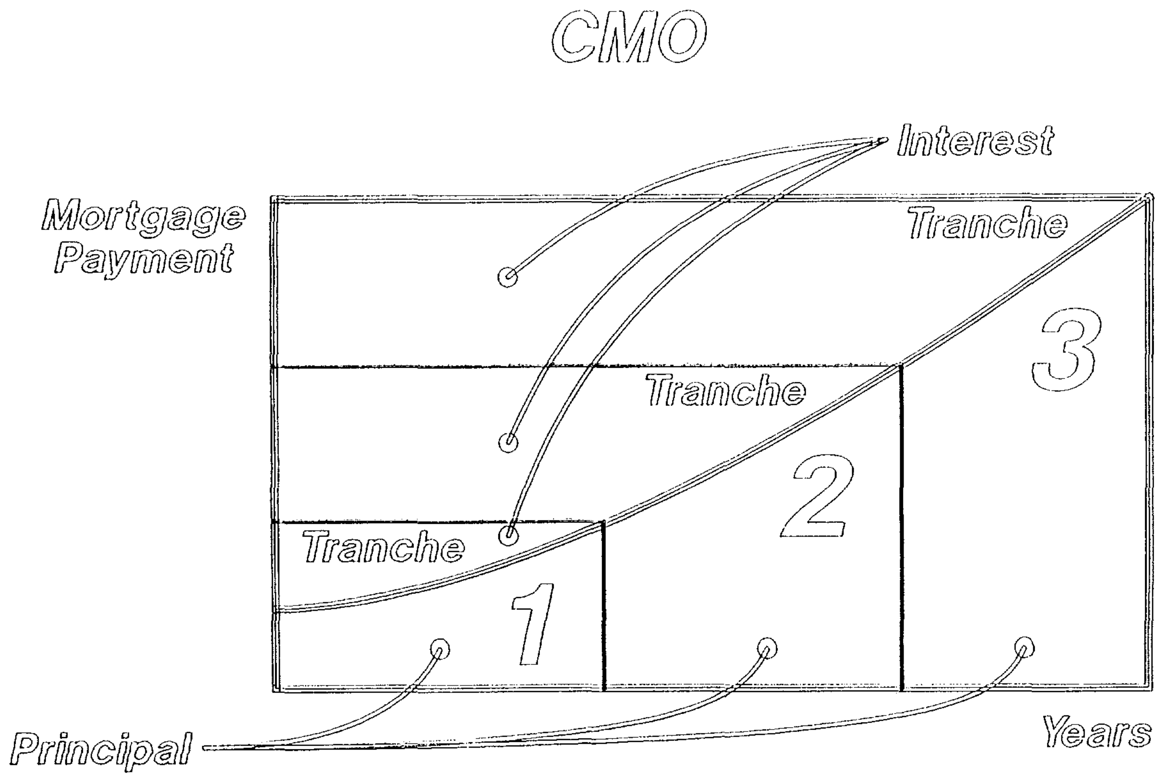
CMO
Interest
Mortgage Payment
Tranche
Tranche
Tranche
1
2
3
Principal
Years

EXHIBIT 4

CMO MECHANICS (WITH Z BOND)

Payment	Tranche 1 Interest	Tranche 1 Principal	Tranche 2 Interest	Tranche 2 Principal	Z Tranche Interest	Z Tranche Principal
1	$34.50	163.80+23.98	$41.62	$	$	$
2	15.62	156.20	41.62	50.36	--	--
3			36.59	227.21	--	--
4			13.89	138.65	--	111.28
5					135.26	128.54
Price	$352.59		$419.98		$227.62	
Yield	8.10%		9.66%		11.23%	
Average Life	1.4 years		2.9 years		4.7 years	

EXHIBIT 5

CMO MECHANICS
WITH PREPAYMENTS ($100 PER YEAR)

	Tranche 1		Tranche 2		Z Tranche	
Payment	Interest	Principal	Interest	Principal	Interest	Principal
1	$34.40	163.80+23.98+100	$41.62	$	$	$
2	5.62	56.20	41.62	260.36	--	--
3			15.59	155.86	--	192.35
4					92.03	47.47
5						
Yield	7.66%		9.56%		11.70%	
Average Life	1.2 years		2.3 years		3.4 years	

EXHIBIT 6
CMO SUMMARY

Tranche	Before Prepayments Yield	Avg. Life	After Prepayments Yield	Avg. Life
1	8.10%	1.4 years	7.66%	1.2 years
2	9.66	2.9	9.56	2.3
Z	11.23	4.7	11.70	3.4

The other interesting thing is, it has a tendency, at least in a normal yield curve environment, to tilt the yields even more severely. The short term reduces its yield but the Z-tranche actually increased its yield with prepayments.

The important thing of CMO mechanics is the characteristics of the underlying asset. You need to understand what you are buying. You need to understand the number of tranches including the Z-tranche. The Z-tranche is extremely important because of the way the interest payments come back to the preceding tranches. You need to know the yield, par, and principal for each of the tranches. And for prepayments, you must project the cash flows of all the preceding tranches, even if you don't own them. So, if your investment person gives you the portfolio of assets with only Z-tranches and you are trying to predict the cash flows under volatile interest environments, you have to predict or project all of the prepayment patterns of the tranches preceding that before you can get to your own asset. It makes for a very complicated mathematical model.

Now going back to the more mundane. In the regular assets (again, these are, I think, fairly obvious) the security detail you need for your bonds, mortgages, etc. -- book, par, and market value; book yield, note rate, and market yield rates; maturity; call/prepayment specifics; quality rating; and marketability. This is all necessary information for you to do your projections.

As for the asset model, there seem to be two schools of thought when you are modeling assets. The one that I like to do the most is seriatim. I don't model assets. I just get a tape or diskette, either from Schedule D or their investment files, and we model each and every asset. Now, when we get into a bigger company that has many assets, some of which can be collapsed down into a model (kind of like the life insurance models that we are used to dealing with), the key parameters, ranked in priority, are asset types (bond, mortgage, etc.), asset features (callable, prepayable, etc.), maturity grouping, and book yield/note rate groupings.

Switching gears a little bit, I want to talk about asset assumptions. There are basically three critical asset assumptions that I would like to address. The first one is "quality spreads" or spreads over treasury. This is, again, the market yield spread over the risk-free treasury. That's how investment people look at the particular yield for a security. The normal form -- the one that our system uses, is a simple mathematical model that says the market yield is equal to A + B times the treasury rate, where the treasury rate is for some particular duration that's associated with the asset that we are dealing with. "A" is a constant, "B" is a multiplier. Where do you get A and B? Two places. You can either use historical analyses or you just walk into your investment department and say, "What should it be?" One comment here: spreads should narrow as assets near maturity. In other words, if you buy a 30-year asset, you may get a spread of 140 points, and you plug it into your little formula and say the market yield is 140 plus one times the treasury. If that is the formula that is locked into that model, as you move through time the spread is always going to be 140 points over treasury. As the treasury slides down, your yield slides down, but it is always going to be locked in at a 140-point spread. I think it is fairly common knowledge in the spread shrinks as you move near maturity. That is true even in an inverted yield environment. The spread on a 20-year instrument is about 85% of what a spread is on a 30-year instrument. The spread on a 10-year instrument is about 85% of what it is on a 20-year instrument. The spread on a 5-year instrument is about 85% of what it is on a 10-year instrument. The spread on a 2-year instrument is about 85% of what it is on a 5-year instrument. So, it is not in equal intervals, but the constant seems to be 85% on kind of a log scale. We have done some modeling that way, and it seems to do a fairly good job. That's extremely important if you are in a buying and selling or a trading of assets scenario because as the maturity shortens and the spread shortens, the market values increase. That becomes very important.

Now I have some graphs of this historical analysis of where I personally come up with the As and the Bs. I think this is publicly available information (Graph 1).

GRAPH 1

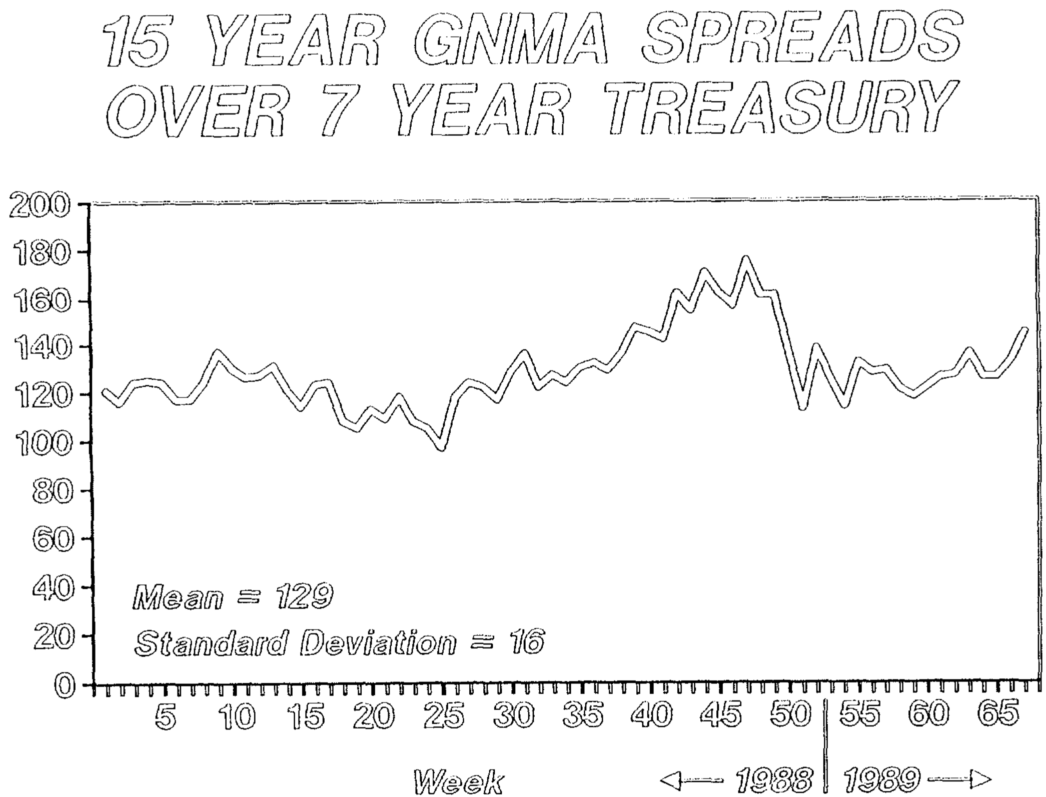
15 YEAR GNMA SPREADS
OVER 7 YEAR TREASURY
200
180
160
140
120
100
80
60
40
20
0
Mean = 129
Standard Deviation = 16
5 10 15 20 25 30 35 40 45 50 55 60 65
Week
1988
1989

What we have done for all the weeks in 1988 and up through mid-April of 1989 is track the spreads of the 15-year Government National Mortgage Association (GNMAs) over 7-year treasury. Over that period of time there is obviously some volatility, but the average spread was 129 basis points and the standard deviation was 16. So with these two numbers you can either create a stochastic model to project what your spreads are or you can use this as a linear regression and come up with an A and a B that will do a good job over a number of interest rates that will give you the 15-year GNMA spreads relative to a 7-year treasury. The reason why I am using a 7-year treasury is the average lifetime for a 15-year GNMA tends to be around 7 years, so I believe the investment community always ties those two together.

On a 30-year GNMA the spread is slightly reduced, and the standard deviation a little bit more (Graph 2). This is over a 10-year treasury. For five-year AA bonds (Graph 3) (these are callable bonds), the mean is 47, the standard deviation is 10. Ten-year AA callable bonds (Graph 4) produce similar numbers -- 60 and a 12 standard deviation. And finally, we have a ten-year BBB is a rating, similar to AA editor bond, callable, where the mean is 109 and the standard deviation is 16 (Graph 5). Again, using this information, you should be able to have a fairly decent model as far as your A and B or your spreads and constant multipliers for getting your market yield.

The next critical assumption is in the prepayment area. There are two prepayment formulas I am used to looking at. One of them is the Public Securities Administration formula, or PSA, which measures the frequency of prepayments on mortgages (Exhibit 7). The PSA model says that the base rate is zero in the first month graded up to 6% (that is an annual prepayment rate) by the 30th month. Then there is some interest sensitivity attached to that, and the prepayment rate is the base rate times the factor. Note the factors for GNMAs and for Federal National Mortgage Association (FNMAs). The column entitled "Note Rate Less Market Rate" is essentially the spread between what your mortgage is at and what you could get in the marketplace. So, common sense says that if you have a 200 point spread, it's time to refinance your mortgage, so that for GNMAs there is a factor of 3.9 and with a base prepayment of 6%, the total prepayment rate is 3.9 x 6 or 23.4%. That's the model that I believe is most often used.

EXHIBIT 7
CRITICAL ASSET ASSUMPTION
PREPAYMENTS

o Public Securities Administration Formula
Base Rate is 0% in first month, graded to 6% by 30th month
Prepayment Rate = Base Rate x Factor

Where:

Note Rate less Market Rate	Factor GNMA	Factor FNMA
+200 bp	3.9	5.1
+150 bp	2.6	3.9
+100 bp	1.6	2.7
+ 50 bp	1.25	2.0
0 bp	1.15	1.6
- 50 bp	1.05	1.45
-100 bp	1.0	1.45
-150 bp	1.0	1.35
-200 bp	.9	1.25

Another model is based on GNMA prepayment information from 1986 through 1988. I looked at the spreads and I came up with a formula of 12% + 4 times the spread + spread^2. That did a very good job of mimicking what the GNMA prepayments were over time. Now, putting this into our computer model and letting the interest volatility affect the spread, I think it's done a good job of representing what prepayments have been.

On the bond side there seem to be two approaches to dealing with calls. I think the most common is a call trigger. Basically, if the difference between the coupon rate and the current market rate exceeds a certain number -- 150, 200 basis points -- then the bond is called. How do you come up with 150 or 200? I have put together a simple example (Exhibit 8).

GRAPH 2

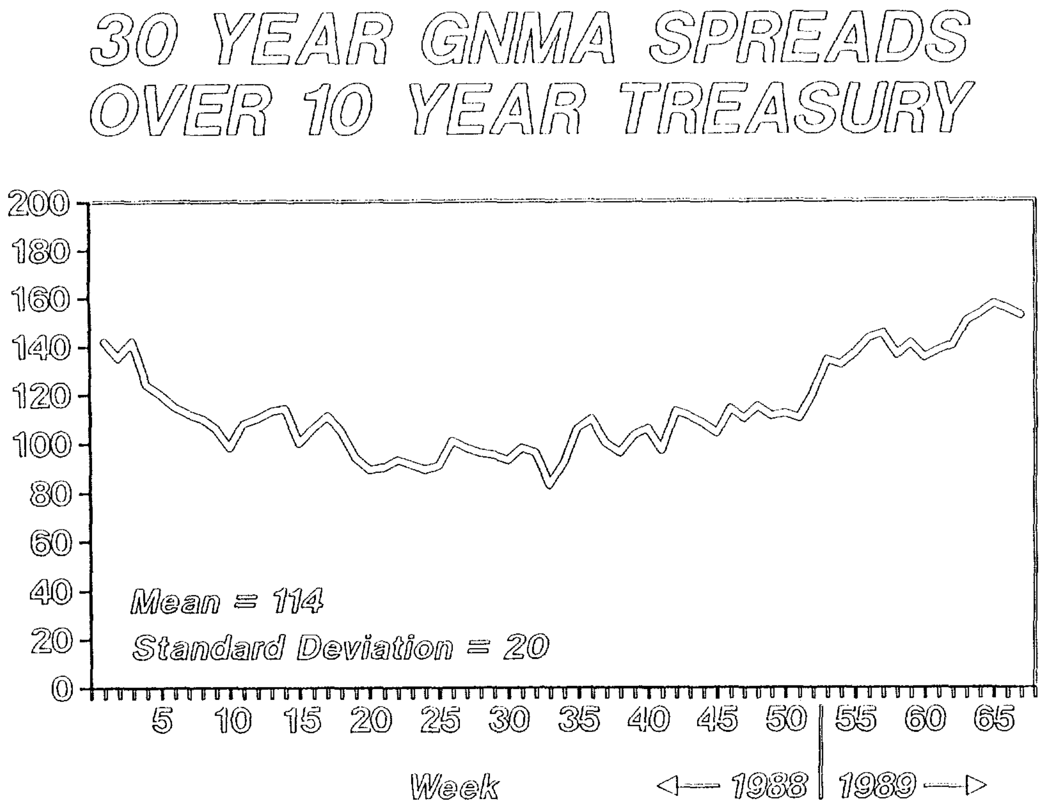
30 YEAR GNMA SPREADS
OVER 10 YEAR TREASURY
200
180
160
140
120
100
80
60
40
20
0
Mean = 114
Standard Deviation = 20
5 10 15 20 25 30 35 40 45 50 55 60 65
Week
1988
1989

5 YEAR AA BOND SPREADS OVER 5 YEAR TREASURY

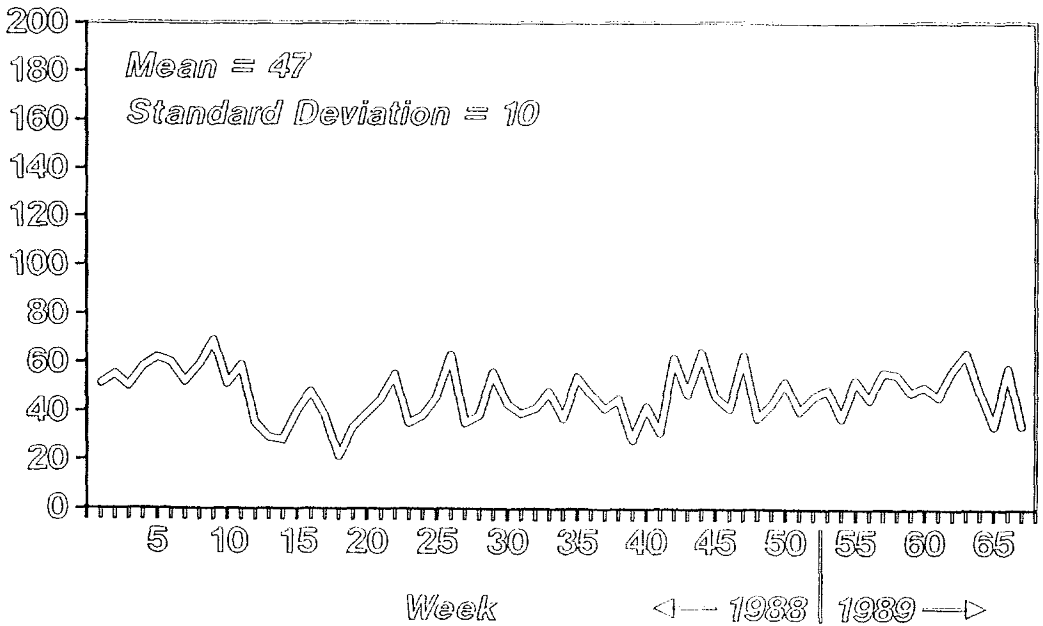

GRAPH 4

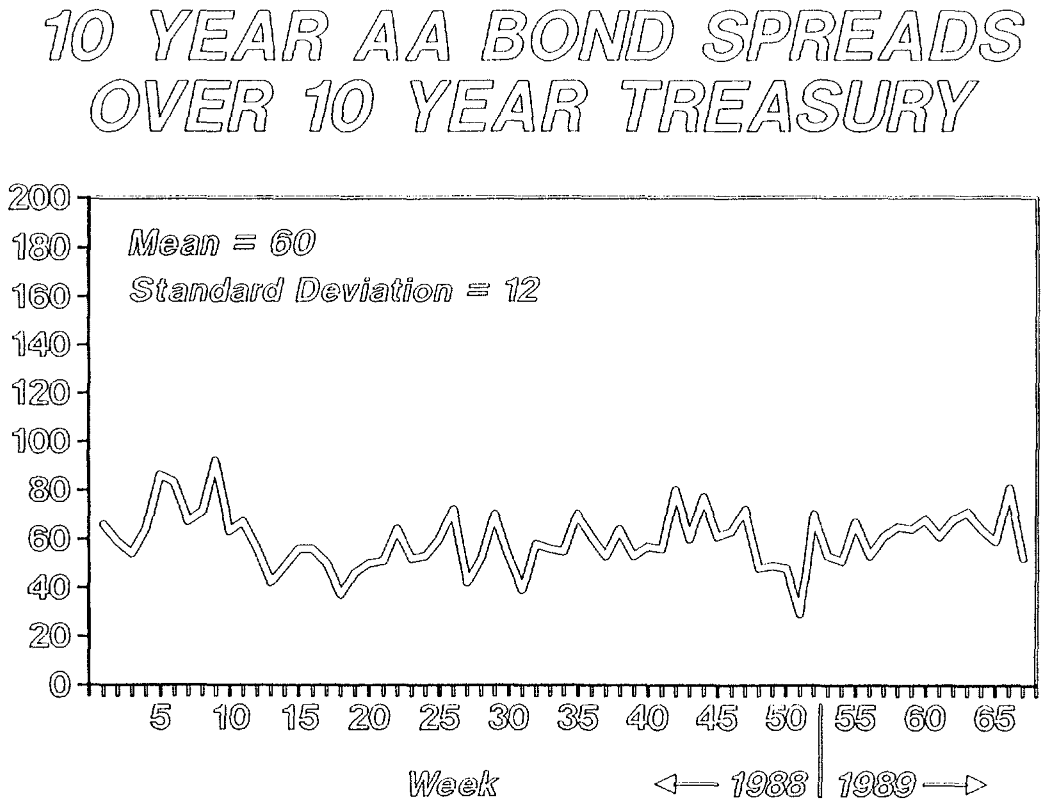
10 YEAR AA BOND SPREADS
OVER 10 YEAR TREASURY
Mean = 60
Standard Deviation = 12
200
180
160
140
120
100
80
60
40
20
0
5
10
15
20
25
30
35
40
45
50
55
60
65
Week
1988
1989

GRAPH 5

10 YEAR BBB BOND SPREADS OVER 10 YEAR TREASURY

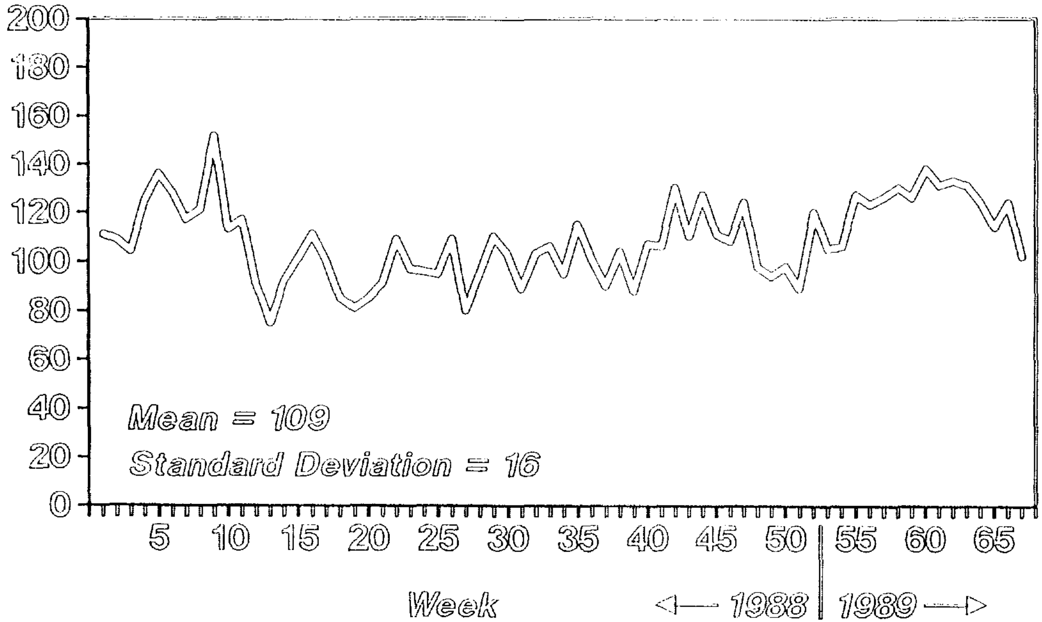

EXHIBIT 8
CRITICAL ASSET ASSUMPTION
CALLS

- Call Trigger
  Bond is called if difference between coupon rate and current market rate exceeds call trigger
  Normally set at 150 or 200 basis points.
- Example: 5 Year Bond, Callable in 3 Years at 102

Coupon Rate	10% Market Yield Price to Maturity	10% Market Yield Price to Call	8% Market Yield Price to Maturity	8% Market Yield Price to Call
6%	84.8	91.5	92.1	96.5
8	92.4	96.5	100.0	101.6
10	100.0	101.5	108.0	106.8
12	107.6	106.4	116.0	111.9
14	115.2	111.5	124.0	117.1

This is a five-year bond callable in 3 years at 102, a fairly typical bond. I have shown different coupon rates and market yield rates of 10% and 8%. Those are the current environment rates. What I have done is show the maturity and the price to call. The boxed-in areas show that it costs less to call than it does to hold it to maturity. Therefore, those are the situations in which it would be prudent to call. If you look at the coupons relative to the yield, in all cases it is 200 basis points. So for this kind of a security, I would stick into my model a 200 basis point call figure.

The other way is, for lack of a better word, economic value. Rather than just making a simple-minded call trigger assumption, why not calculate for each security present value of future cash flows with the call at the various interest rates you're under and see if it makes economic sense to call it or not. That requires a lot more computing time. It's not a simple "compare spread versus two" and call it or not call it, but with current technology and with computer power, I believe we will see most computer models heading this way as far as handling calls.

The last critical asset assumption that I will talk about briefly is defaults. There are two ways of handling those. The first way I will call the "hold back" method. It works this way: When you model you don't take into account Mandatory Securities Validation Reserve (MSVR). In your projections of asset cash flows you reduce your investment income by a certain charge to investment income to cover defaults. The idea is that it is an expense. It leaves your asset model and goes over into MSVR. These are normal holdbacks that I guess I have seen and I have used and feel comfortable with, a 150 basis point holdback for a non-investment grade, 10 basis points for an investment grade, and 30 basis points for a commercial mortgage or a private placement. There obviously is no hold back for GNMAs or agency issues or treasuries.

The other point I want to make was promulgated by the New York Insurance Department, and as a fallback a lot of companies use this because they don't have to think about it. New York Regulation 126 states that if you are not going to be able to come up with holdbacks with which your Investment Department feels comfortable, you have to use at least 75% of your MSVR charge. That's an easy fallback. So some companies have used that. The problem is that the MSVR misses a lot of assets like mortgages and private placements. I believe that the holdbacks that I have shown here are probably more prudent.

The second method of handling defaults is the stochastic method. This involves estimating asset defaults based on historical default information. Using this method, MSVR is included in the asset base, and investment income is not reduced for any holdback. The stochastic model projects some asset defaults which directly impact the cash flows.

The tables of historical default rates are generally available numbers (Exhibit 9). The 1980 to 1988 estimate is 25 basis points with a standard deviation of 36. That's from some information I was able to obtain recently.

EXHIBIT 9

CRITICAL ASSET ASSUMPTIONS
DEFAULTS (CONTINUED)

- o Stochastic Model
  - MSVR
  - No explicit reduction in investment income
  - Project asset defaults based on economic scenarios
  - Historical default rates:

Period	Mean	Standard Dev.
1900-1909	.89%	.67%
1910-1919	2.01	1.24
1920-1929	.95	.40
1930-1939	3.20	1.84
1940-1949	.43	.55
1950-1959	.04	.05
1960-1969	.03	.04
1970-1979	.13	.18
1980-1988 (est.)	.25	.36

If you take this sort of distribution of the mean and the standard deviation and you use some of your risk theory, you can probably create a stochastic model with the distribution that produces default rates similar to this under certain economic conditions. You use that in your asset model and you can project defaults into the future. It is important that you handle defaults one way or the other. This is probably the more scientific way, but because there are a lot of assumptions that people probably don't feel comfortable with yet I think the more common approach is what I showed previously, which is the holdback method.

The last thing I want to talk about is asset strategies. It seems there are two areas that are important when you talk about asset strategies -- managing existing assets and investing new cash flows. Managing existing assets covers two issues. One is a buy-and-hold. What are you going to do with the assets you have on the books? A buy-and-hold is not an interest-sensitive strategy. You just sit on those assets and let them go. When I say you sit on them, if calls and prepayments are affected, then the model should reflect that, but you do not consciously disinvest them. That is a common strategy. It is important that you talk to your investment people. Otherwise, the results of your projection, if you are assuming a buy and hold and they are trading 75% of their portfolio, show we are not in the same world. Nobody is going to believe the numbers.

The other is an active strategy which I'll call interest sensitive. It ought to be based on where interest rates are. I will quote from some of our client work a disinvestment strategy for existing assets. "An active strategy is followed for mortgage-backed securities because of the high liquidity and efficiency of the marketplace. For all other assets -- bonds, private placements, and so on -- a more passive or buy and hold strategy is followed. The trading strategy for mortgage-backed securities depends on the perceived interest rate environment. It is summarized as follows: under a normal or positive yield curve where interest rates are stable, there is no trading because there don't seem to be significant opportunities in the marketplace. Under a normal positive yield curve with interest rates rising, active trading should take place to capitalize on the opportunities in the marketplace and to further shorten the portfolio. Assume that the active trading occurs whenever a capital gain of 16/32 of 1% of par or more can be achieved, assume that the trading costs total 4/32 of 1% of par. In a normal positive yield curve environment where interest rates are falling, an active strategy would take place to capitalize on opportunities and further lengthen the portfolio. Assume that active trading occurs whenever a capital gain of 16/32% of par or more can be achieved and the trading costs are also 4/32%. And under a negative or inverted yield curve environment there is no trading that takes place." That was not an actuarial dissertation in investment strategy. That was me walking into the investment department of the client and explaining what we are trying to do with our process and having him write it down on a piece of paper. The minute he wrote it down on a piece of paper, we were able to stick it into a computer model and emulate what was going on.

I challenge you. That's what you need to do too. You have to do that; otherwise, the credibility of your results is going to be woefully inadequate and not accepted at all.

That is what I am talking about as far as an active investment strategy for existing assets. As far as investing new cash flows, there are two issues involved. Sometimes there are negative cash flows, in which case you can do one of two things. You can borrow or you can liquidate. You need to model that. And finally, if you have positive cash flows, you can do the famous "old case study life" approach and everything goes into 5-year callable bonds. Then there is the famous dynamic strategy, similar to what I described here. You sit down with your investment officer and get in writing what an investment strategy is under volatile interest environments and try to create a computer model that emulates that. We are not trying to take over his job. We are trying to see what he does and what that does to cash flows under volatile interest environments.

MR. SHELDON EPSTEIN: I'm going to talk about, in particular, Merrill Lynch's new approach or evolutionary approach to asset/liability management. And what it entails is a synthesis of some of the buzz words that we've been hearing about for the last couple of years . . . option pricing, simulation, and something that's new, yield curve dynamics. This last portion is what I want to focus on because I'm most excited about this and its implications for defining investment strategies that will fund liability targets.

Basically, I want to start out with our philosophy here. Our philosophy for either an asset or a liability is to determine its market value under various scenarios. This assumes, at this point, that we have a black box that says you can calculate what the market value of your CMO is, if the yield curve is upwardly sloping or downwardly sloping or parallel shifting up or down. If we have that, then we can develop return and price sensitivity measures for the security or the liability again, based on starting systematic yield curve dynamics. How does a yield curve actually move? And then we come up with a set of summary statistics for each security and let the investment manager determine whether it's appropriate to include that security in his portfolio, given that he has a certain type of liability that he has to fund. Finally, I want to briefly go into the idea behind option pricing models and take some of the mystery out of option pricing models.

Wall Street has traditionally presented these as mysterious black boxes that do all these wonderful things. Well, they're basically just sophisticated interest rate path generators. The models that I'm acquainted with are very similar to some of the models that Joe and Greg and some of the other actuarial consultants are using. The only difference is that the interest rate generator has to satisfy three important conditions. First, it must price all non-interest sensitive cash flows correctly. Remember, we're coming up with a pricing mechanism and if we have a pricing mechanism that doesn't price, for example, a treasury bond correctly, there's something wrong with the pricing mechanism. But basically what we're doing is generating interest rate paths and determining the cash flows along those paths and somehow discounting them back to come up with unique prices. Another important condition is that there are no negative interest rates. This requirement is necessary so that arbitrages don't exist. If you knew that there are no negative interest rate and you could sell a bond with a negative interest rate, you could sell the bond, get the cash and pay the guy back less money when the bond matures and come up with a risk-free profit. Finally, the interest rate paths themselves have to demonstrate the volatility that you wanted. A number of option pricing models I've seen will satisfy one or two at the expense of three. What we're trying to do is come up with a model that will have the volatility that we want, that satisfies all the pricing constraints. By the way, if you satisfy these three conditions, you also satisfy what's known as put/call parity. Wall Street, again, has tried to say that that's the condition that you're trying to satisfy. That's a result of satisfying these three conditions and it's something you don't really have to get into in your computer modeling of the interest rate paths themselves.

Well, the whole idea of generating interest rate paths is you have to have some way of moving yield curves through time. In the language that I'm familiar with, we call this a diffusion process and it refers generally to the probability distribution that's applied to the initial yield curve. Common diffusion processes are log normal. Joe has talked about that type of diffusion process before, and basically it's a normal distribution for the interest rates. You have to put constraints about mean reverting on it so that you don't have negative interest rates. The important thing is that if you have interest rate paths that demonstrate a certain type of volatility, the volatility you want to show, the choice of diffusion process is relatively immaterial. It's material if you're talking about a few basis point differences in yield, but in terms of getting the aggregate characteristics of the security, if you're comfortable with one type of diffusion process or another, it doesn't really make that much difference. However, I'm going to counter that by saying that some of the research that we're showing and some of the results I'm going to show you are going to come up with a diffusion process that will allow you to describe full yield curve dynamics with relatively few parameters. That's going to make it relatively practical for people to start actually using these things to analyze securities on the run.

Simulation models are what we're all familiar with. The option pricing paths are used as the interest rate paths for the cash flow generator. In particular, I use one of the actuarial simulation models as applied to the interest rate paths that I generate with my option pricing model for insurance liabilities. For assets, we have special cash flow models for mortgage-backed securities, for CMOs, for options and futures, you name it. We have a cash flow generator, but the basic idea is that the cash flow generator uses the option pricing paths as input. The cash flows that we focus on are only the actual cash movements, the real cash flows. Statutory and GAAP flows can be used to simulate the results for accounting purposes, but in terms of pricing, in terms of coming up with prices that the capital markets would come up with, you can ignore the statutory and GAAP flows. In summary, the option pricing mechanism is a sum of option pricing paths and a cash flow generator. For the rest of my talk, I'm just going to assume that we have this black box that says if you put this yield curve in and if you put a certain volatility in, you'll come up with the price. The cash flow assumptions are tricky, but they can be assumed.

For example, in modeling a single premium deferred annuity (SPDA), one of the points of contention is always the lapse function. Similarly, in the mortgage-backed role, it's the prepayment function. However, you always are making some kind of implicit assumption in whatever modeling you're doing for these parameters and the point is you can test how the price of the security will change as you change that particular assumption.

Now I'm going to go back in time to when I was writing actuarial exams and I'm sure everybody's familiar with Reddington's immunization paper. All that we are trying to do on Wall Street is validate his theory, and his theory was that if you could match up the price sensitivity of your assets and your liabilities, and if you could set up your asset portfolio in such a way that a number which is called convexity would be greater than the convexity of the liabilities, then no matter how interest rates moved, you would come out ahead in terms of market value and that's really what counts.

Duration is the first derivative of price with respect to a change in interest rates and convexity is the second derivative. One of the points about that though is that when we apply it now, we're talking about an active or what we're going to call an interest-sensitive strategy; that is, the duration targets, the convexity targets, change over time, so the only way to keep that type of structure in force is to actively re-balance. However, the re-balancing tends to be minimal and I'll give you an example from real life at the end of my talk. The point about Reddington immunization was that he originally developed his formulas for duration based on non-interest sensitive cash flows and parallel yield curve shifts. The di in that derivative could be any type of interest rate shift. It just so happens that you get nice formulas when you have non-interest sensitive cash flows and parallel yield curve shifts. We apply the same formula, or Wall Street does in general, when they use option pricing. The question is: how do you shift the yield curve? And what has become the standard is to shift it in a parallel manner. What we've done is we've gone back to first principles and we've said "It seems there's a lot of evidence that yield curves don't necessarily move in a parallel fashion." And one of the reasons why the strategies that we would normally recommend would break down is precisely that we're ignoring a vast range of types of yield curve movements.

I'm going to give an example of how this breaks down. The point is that duration and convexity can be measured for securities based on parallel shifts, but I'm going to show you duration and convexity measures based on other types of shifts that will determine immunizing strategies that are more robust, number one, but number two, that will give you a basis for taking a risk position so that you know exactly what the level of risk you are taking is. Given that, being perfectly hedged does cost something. It's not generally free.

The other point is that you can define duration and convexities on a forward basis. If you want to manage so that you only re-balance once every six months, you can calculate forward durations.

What will the duration of my assets be six months from now? What will the convexity of my assets be six months from now? That way, as long as you have the ability to withstand market value changes in between the present time and your time horizon, you will make sure that, at the time you want to re-balance, your assets and liabilities will be in the correct balance.

There are a lot of confusing descriptions about convexity. Basically, Reddington immunization says we have a price curve for the liabilities. We want to create a price curve for the assets, which is the yellow line. We want the assets to be worth more than the liabilities if interest rates go down, for example, in a parallel shift of 100 basis points, or if interest rates go up with a parallel shift of 100 basis points. Duration is calculated as a measure of the slope of the price curve. If we have durations equal, for small interest rate changes we don't have to worry about anything else. The price of the assets will track the price of the liabilities. However, for large interest rate changes, because of the curvature of the price curves in general, you'll get some error if you just use the straight line approach. That's where convexity comes in. Convexity is just a measure of how far the price curve deviates from the tangent line. The higher the convexity, in general, the higher the curve will be above the tangent line with larger and larger interest rate movements. So, if the assets have a higher convexity, because they deviate more both with upward parallel shifts and downward parallel shifts in interest rates than the liabilities do, we could be happy that we've hedged out any interest rate risk, assuming that interest rates just move in a parallel fashion.

However, as I mentioned, the interest rates don't just move in a parallel fashion and this example will point that out. I'm going to assume right now that the insurance company can issue a five-year $100 bond at the treasury curve levels. That is, they can issue it with the same coupon and yields from, in this case, a June 30, 1988 U.S. Treasury curve. I'm going to assume that they're going to invest the proceeds in some combination of two, 10- and 30 year-bonds that will immunize them from interest rate risk. That may not be the strategy they want to take, but in order to determine the strategy they want to take, if they're bearish on interest rates or bullish, we have to at least determine what is the risk neutral or, if you think of it in the efficient frontier mode, what is the zero risk reward portfolio.

To give you an idea of what the yield curve looked like on June 30th, it was upwardly sloping. Well, it's actually gone back to an upwardly sloping. The 90-day rates were fairly low. They were about 6.70%. For our particular bond that we're issuing, which is a five-year bond, the coupon would be 8.40%. The two-year bond, which we would look at as an investment opportunity, had a yield of 7.98% and a coupon of 7.98%. Again, we're going to be investing at current coupon bonds. A 10-year bond had a coupon of 8.81% and a 30-year bond had a coupon of 8.93%. Again, Reddington's formula or strategy for immunization is to come up with a portfolio that has exactly the same duration as the liability, but which has the most convexity. And what always happens when you stick this into an optimization model with those constraints, you want to maximize convexity and minimize price, is that you end up with what's called a barbell strategy and that's what happened here. Basically, it tells us that in order to get the highest convexity with those three securities, we would invest $74.48 of the $100 of proceeds that we got in the two-year bond and $25.52 into the 30-year bond. If that were the case, we would then have a duration of four years for the assets and four years for our liability and the convexity of the assets would be much greater than the convexity of the liabilities.

Well, I want to test what happens when a yield curve moves. First, I'm going to start out with the June 30th treasury curve, and I'm going to parallel shift it up by 50 basis points and I'm going to see what happens to the price of my assets versus the price of my liability in that instantaneous type of parallel shift. Then I'm going to look at it with another type of yield curve movement. I'm going to say the yield curve tilts, to become steeper. I call this a negative tilt and I'll explain why later, but I want to see what the immunization result was under that type of yield curve movement, which is not unlikely to occur. The point is that under the original portfolio, if there wasn't a change and we're just looking at instantaneous changes in interest rates right now, the value of the assets and the values of the liabilities are equal, as they have to be. If interest rates move up, the values of the assets go down, as do the value of the liabilities, but the liabilities go down more than the assets because of the fact the assets have higher convexity. That's what we wanted. That is the result we tried to achieve with our immunization strategy and we may have forced that profit.

However, when the yield curve tilted, the two-year bond went up in value, the 30-year bond went down in value and the five-year bond stayed at about the same level. It went down a little bit. For the net result we would have had an economic loss, a market value loss in our net portfolio of eighty cents. Obviously, immunization, as stated now and as most Wall Street people try to communicate it, will break down with anything other than a parallel interest rate movement. What we had to do, and this is what I'm most excited about, is go back to first principles.

How do yield curves actually move? We have a group at Merrill Lynch called the Financial Strategies Group, which includes a lot of the so-called rocket scientists. In fact, one of the so-called rocket scientists really worked for NASA. That's Dr. Herman. Dr. Dash, and he's not the evil Dr. Dash, and Dr. Balas are two people who have developed this idea. Jerry Herman actually wrote the first paper on this in 1987 and it's now where we're starting to incorporate this into the option pricing models that we're using. It is a method to decompose yield curve movements into independent yield shift modes, so into a parallel shift mode, into a tilt mode, into a bend mode, into S-shaped yield curve changes, and to try and come up with yield curve movements which are independent of each other. There's a mathematical technique they use. It's called empirical orthogonal functions. If you think back to part one where you studied eigenvectors, it's basically coming up with vectors of yield curve movements, vectors applied to the whole yield curve that are independent of each other that, in a linear combination, will produce the yield curve movements from today to tomorrow, for example. It's not important how they actually do it, but the result is actually quite startling.

What they do is come up with a matrix which basically shows the yield curve changes on a day-to-day basis or a month-to-month basis, and in most cases daily, along the maturity structure of the yield curve. So they might calculate for the first check date that the 90-day rate went up by four basis points, and these could be the log normal changes or they could be the absolute changes in yield. And at the other end, the 30-year bond rate fell by three basis points. This would be a case where the yield curve inverted a little bit. It's not quite clear. You end up with a matrix of yield curve changes. Looking at a number of specific maturities seems to be the standard way to present a yield curve. For example, if I was looking at 11 maturities, I could come up with 11 vectors of yield curve movements that will totally define, in any linear combination, all possible yield curve movements. For example, if you had a vector that said that the 90-day rate increased by 100 basis points and everything else stayed the same and another vector that showed that the 6-month rate increased by 100 basis points but all other yields stayed the same, you could obviously come up with a linear combination of those types of vectors that would look like any of the actual yield curve movements we had. The amazing thing is that the first three vectors that we calculate, when we do this type of analysis, explain more than 97% of all yield curve movements no matter how many points, how dense our maturity structure is on the yield curve. These three movements are parallel shift, tilt and that applies to the slope of the curve, and something we call flex or the curvature of the curve.

The other eight vectors that we would have calculated using this type of analysis would only account for less than 3% of total yield curve movements over time. Actually, over the last year, those three movements will account for about 98%, 99% of the yield curve movements. It's interesting that when you do this type of analysis you find out or you can actually measure the contribution of each of these types of movements to overall movements. For example, over the last year, parallel shifts accounted for 40% of yield curve movements; tilts accounted for 40% and flexes accounted for 20% of yield curve movements. How do we use this and what is its purpose?

Well, standard convexity is a second-order measure of parallel shift duration and everybody's focusing on convexity and buying convexity cheap, making sure you have lots of positive convexity in your portfolio. What we feel that people should be focusing on is the first-order measures of all the major yield curve dynamics. Since there's only three of them, it means looking at three duration type numbers. You can then layer on convexity numbers for each of those types of movements. What we do is we actually define a parallel shift duration that's going to be equivalent to the option-adjusted duration measure which Wall Street normally calculates, and also tilt durations and flex durations for changes in what I'm calling normalized parallel shift, tilt and flex vectors.

The vectors that we would have derived from parallel shift, tilt and flex over the period from May 1988 to March 1989 are based on daily yield curve movements. The parallel shift vector is a straight line. Now, we've normalized it so that the parallel shift vector is at 100 basis points and we've taken the actual vector for the tilt and multiplied that by the appropriate factors to get it up to the equivalent of 100 basis point parallel shifts. So when we talk about shifting the yield curve with a parallel shift of 100 basis points or tilting it by 100 basis points, we're talking about adding a vector to the current yield curve. The tilt vector seems to have a pivot point of around four years. In fact, we've looked over the last ten years and that pivot point has remained fairly stable. It's been between four and five years.

I'm going to calculate a tilt duration. I'm going to start out with my original yield curve, calculate a price. Then I'm going to add, for example with a 50 basis point tilt shift, half of the numbers that apply to the tilt curve to the yield curve, which will imply an inversion in the yield curve, and when I have a negative tilt in the yield curve, I will multiply the numbers by minus .5 for the tilt curve. And doing that, you can come up with a price series which defines the price curve, given tilts in the yield curve. How does this work? Well, for example, those vectors were calculated based on daily movements from May 1988 to March 1989. I just wanted to show you what actually happened between June 30 and December 31. We started with the June 30 treasury curve. The arrows on this graph are incorrect. What we actually had was an 88 basis point parallel shift up. Then we had a tilt and that was to the sort of beige line. Then we had another tilt that was about 45 basis points; the equivalent of multiplying that tilt vector by .45. Then we had a flex or a positive bending of the curve to the curve of 106 basis points. We multiplied that flex vector by 1.06. The actual December 31, 1988 treasury curve total of the adjustments I have described. What this means, is that we've been able to define the total dynamics that moved the yield curve from the one curve to another, the June 30 curve to the December curve, with only three types of movements. And you can replicate this over any historical time horizon. It's actually quite astounding. Dr. Dash gets mad at me when I refer to this as cold fusion in a jar, but he prefers super-conductivity, because he's a physicist.

I'm going to go back to our example where we issued a five-year bond and we tried to hedge out all the interest rate risk. What I'm going to do is I'm going to have to invest in at least four bonds to match each of the parallel durations, the tilt durations, and the flex durations, but I also wanted to have some positive convexity on a parallel mode, so I need actually five bonds in order to match all those measures of sensitivity. We ran this through the optimizer. It's a linear program. It's basically pretty simple to calculate and you find out the proportions of each bond you should buy in order to match the price behavior of the five-year bond. For example, here we had to go short the one-year bond, go long the two year, buy more than the amount of bonds that we actually sold, go long a little bit on the seven year, go long on the ten year and go short the 30 year. What this is telling us is it's pretty expensive to hedge this particular security if you can't buy the exactly offsetting security and so you would probably, in reality, want to hedge it with some type of strategy that didn't perfectly match all of these durations, but you would base the amount of risk you were willing to take in each of the dimensions off of this sort of risk neutral table. Well, what happens when we change the yield curve if we actually invested in that portfolio? What I did, and one thing to point out is that those durations mentioned before were based on forward durations. That is, how would the prices change six months down the road, because I assumed that the insurance company was not going to re-balance for six months. That's why the five-year bond had a parallel duration of 3.7 years versus four years when we were looking at instantaneous shocks. In this instance, I just parallel shifted up the June 30 curve by 50 basis points and assumed that was the curve on December 31. In that instance, you actually break even on the performance of the assets versus the liabilities. With a minus 50 basis points tilt, that's the yield curve becoming steeper, you lose a couple of cents, but it's nowhere near the eighty cents you would have lost if you had just used traditional immunization techniques. With a flexing of 50 basis points, you would have actually broken even again. Now the interesting question is what happens if you actually plug in the December 31 yield curve. You end up with a gain of 13 cents. The reason you have such a large gain is because I forced it to have positive convexity in the parallel shift mode. We had an 88 basis point parallel shift in the parallel mode between June 30 and December 31, so you get a lot of pick-up. We didn't re-balance over a large interest rate change.

In summary, what we're advocating is a systematic analysis of price behavior over the widest range of independent scenarios and I can't stress it enough. Each of those vectors is independent. They're not correlated in any way. So if we can just analyze the price of our security, whether it's an asset or a liability, in five parallel shift modes, in five tilt modes, in five flex modes, we're going to insure that the assets and the liabilities perform as desired, because any other yield curve movement is just a linear combination of those three types of movements.

What I'm currently working on with all of our high-powered physicists and meteorologists, etc., is implementing this inside the option pricing model as a diffusion process. What it means is that traditionally, in order to get an option pricing model to work with say a log normal process, you had to define this humongous covariance matrix between the various maturities of the yield curve to find the correlation between say the 90-day rate and the three-year rate and the 90-day rate and the 30-year rate. With this, you just have to define three types of volatilities: the volatility of parallel shifts, the volatility of tilts and the volatility of flexes, and you actually get as a result the focal variance matrix that you would have had if you used one of these more complicated models. What it means is that the computational time to produce an option pricing model which is arbitrage free has dropped in the magnitude of about 500%, and what it also means is now we can systematically determine how the prices of options embedded in securities will change if volatilities change in a parallel mode, in a tilt mode, and in a flex mode. This is not the standard. This is real state-of-the-art and we're just implementing it now, but these are some of the things I'm implementing in the liability analysis and my friends on the asset side at Merrill are implementing there. I just wanted to share that with you.

MR. EDWARD C. BYRD: Sheldon, could you expand a little bit on what happens when we generate the yield curves?

MR. EPSTEIN: Are you talking about the option pricing model?

MR. BYRD: Yes.

MR. EPSTEIN: Okay. What is basically happening is we're generating interest rate paths. In particular, I prefer to use, and this is personal preference, a log normal type of diffusion for interest rates, so to generate each path I shock the interest up or down, by random numbers that I sample from a log normal process. The main thing that I'm doing differently than what, for example, Joe does in his modeling is I don't assume that the mean of that distribution is zero. I assume there's some other mean because I have to satisfy a condition that when I price a cash flow that's certain, I come up with the price that I would have come up with if I had used traditional bond mathematics. Doing that will make your model arbitrage free. When I do that, I then have a bunch of interest rate paths that I know represent a certain volatility and I now price non-interest sensitive cash flows correctly. I can then plug that into an actuarial simulation model to generate the cash flows from an insurance product, for example, or to generate cash flows from an asset if I'm looking at a mortgage-backed security. Then what happens is I do that for 100 or 200 or 300 different interest rate paths. I'll generate it off of an initial yield curve with an initial volatility assumption and I discount those cash flows along each path's interest rates, short-term interest rates, come up with an average price, an average present value and that is the price. I know it creates prices correctly for non-interest sensitive cash flows. I can test how well it produces prices for things like callable bonds or mortgage-backed securities by looking in the marketplace. I can call the traders downstairs and find out what they're bidding on those types of securities and, as a result, when I apply it to something that's not necessarily traded in the secondary market, I can be pretty sure that applying the same approach to a different set of cash flows will create a price that looks like something that secondary markets might put on the particular security. Does that answer your question?

MR. BYRD: I was wondering more in terms of maybe something like a particular asset, maybe like an adjustable rate mortgage (ARM) or something like that. What would go into that?

MR. EPSTEIN: Well, for a mortgage-backed security, we would have a pre-payment function that says that the amounts of mortgages that would be pre-paid in any particular time frame will be based on interest rates at that horizon and some history of interest rates and we make similar types of assumptions on the insurance side. We'd assume that homeowners aren't deficient and they have some type of memory when they decide whether to pre-pay their mortgages or not and there will be a base pre-payment function very similar to what Greg was talking about. If I have an interest rate path that gives me the full treasury yield curve and I assume that my pre-payments are based off of the treasury yield curve, the function of the treasury yield curve and the history of treasury yield curves in that particular path, I can develop, just using mortgage mathematics, what the amount of outstanding principal will be at any point on that path and what the actual cash flows would be along that path. I also have to build in the way that the rate will re-set on the ARM and that could be a function of Treasuries or some other yield curve that I would generate in tandem, if it's cost of funds or based floating off of treasuries; somehow generate what the mortgage rate will be period by period.

Once I have that, I have a full set of cash flows along that interest rate path for that particular ARM. We then discount back those cash flows using 90-day rates and start out at the tail end of your projection, usually when the ARM expires, as it matures, and discount those back and add in new cash flows as you go through time using the short-term rate; actually one-month rates for an ARM model. You then would have a present value along that path for that particular ARM cash flow. That's not the present value. That's not the market value of that ARM. We have to do that on a hundred different paths and that, in aggregate, satisfies all these pricing constraints that we put on the option pricing model. If we then do that, average the prices that come as a result of each of those 100 paths, we will come out with a number that should replicate the market value of the security. Now it generally won't because the securities can trade at some spread over treasuries, so what we have to do is iterate to an initial yield curve which has a spread over treasuries that will reproduce the market price. Once we have that, we can then shock that yield curve, which is the treasury curve plus what is called the option-adjusted spread, and generate 500 new interest rate paths to determine what happens to the price of this ARM when interest rates have a parallel shift up, for example, and that's how we would calculate durations and convexities. But the general technique is pretty straightforward.

MR. BUFF: Are there any other questions?

MR. STEPHEN A. J. SEDLAK: I was wondering what the equivalent of tilt and flex are on the liability side of product.

MR. EPSTEIN: Well, basically what we do on the liability side is we price the liabilities, again assuming an initial yield curve. We use the same option pricing models and come up with interest rate paths that are plugged into the actuarial simulation model; come up with the pure cash flows, not the statatory cash flows or the GAAP cash flows, though we can do projections to see what happens if you put a particular investment strategy on, which might be the duration match. So what we do is, again, start out with an initial yield curve and we say, for example, if we're pricing an SPDA, we're going to get $1,000 of premium. We need the present value of this SPDA on a market value basis to equal $1,000 and you iterate towards some split that will create a price of $1,000 and that's the spread, the funding target spread for the liabilities. We then shift that initial treasuries plus spread curve, either with a tilt or a flex or a parallel shift, and come up with a new price for the liability. That says, what will the price of the liabilities look like under this different type of yield curve? It could be six months forward, etc. We know we've got $1,000 in cash. We have to invest that in some kind of security which will yield the spread that the liability requires and we have to make sure that the value of that asset moves in conjunction in each of those different types of shift environments in line with the liabilities. If that were the case, we would be sure that we would always have enough assets to meet the liabilities and we'd be needing the yield or the total return of the liabilities. We don't usually think of it that way, but the total return or the yield of the liabilities would be net.

MR. ALLAN MING FEN: One of the advantages you mentioned, Sheldon, with your yield curve dynamics and multi-factoring was the independence of the three different types of factor . . . the parallel, the tilt and the flex. How did you come to conclude that those three were independent?

MR. EPSTEIN: This is going to be a somewhat technical answer. If you remember in linear algebra, if you have vectors in say a three-dimensional space, you can define any other three points. Three points define a point in space. You can define that as a linear combination of three base vectors -- one, zero, zero -- zero, one, zero -- zero, zero, one -- which are all independent. You can have a transformation where you transform the axis and you can define three different vectors which are all independent, they're orthogonal to each other, which can still define that point. The linear combination of those three vectors defines that point. What we've done is created a tilt vector, which is orthogonal to the parallel shift vector and a flex vector which is orthogonal to both the parallel shift and the tilt vector. In effect, you could not produce the yield shift movement with any two of those or any one of those; you need three; in general, three of those to produce the actual shift that would occur. Now the point is if you have 11 points in your maturity structure, which is what I showed you, in general you could always construct 11 vectors that will do that, but you can actually condense it to three of these which are independent. So just generating three random numbers and applying those random numbers with a certain volatility to each of those vectors will produce yield curve movements which are consistent with the way that yield curves actually move. Some of the things that our Financial Strategies Group have done is calculate advanced kinds of statistics to measure how closely our interest rate paths represent the statistics of actual interest rate paths. If you use a normal distribution, you don't get very good, you fail the test very quickly. If you use a log normal, you fail the test a little bit later. If you use this more complicated, but it's actually simple, type of diffusion process, you can replicate the characteristics of actual interest rate paths, the last ten years of data or the last five years or various periods of data, almost exactly. So that's the justification as to why these things are independent and why we can claim that if you analyze the price behavior along each of the axes, any other price movement is going to be a result of three independent movements. So you should be okay with a combination of those movements.