RECORD OF SOCIETY OF ACTUARIES
1993 VOL. 19 NO. 4B

MEASURING COLLATERALIZED MORTGAGE OBLIGATION (CMO) CASH-FLOW VARIABILITY: REGULATORY DEVELOPMENTS

Moderator: DAVID A. HALL
Panelists: CHRISTOPHER T. ANDERSON*
ANDREW S. DAVIDSON†
MIKE H. SIEGEL‡
Recorder: DAVID A. HALL

Insurance regulators have commissioned a task force to develop cash flow variability measurement standards. This session will present the recommendations of that task force.

- Review of objectives
- Measurement standards
- Variability versus risk
- Relationship to cash-flow testing

MR. DAVID A. HALL: I'm senior vice president and actuary for the ITT Hartford Life Companies, where I am currently the director of invested assets. I have also served for the past year as chairperson of the Investment Section, which is sponsoring this session.

Chris has been active in insurance regulation for the last several years. As a member of the fixed-income research group at Merrill Lynch, he works full time on insurance-industry investment matters. Within the NAIC structure, he is active in issues related to reserves and model investment law, and is part of the Invested and Admitted Asset Working Group. For the last year, in this capacity, Chris has chaired the CMO cash-flow variability team. Chris is a graduate of Brown University and received his MBA at the University of Chicago. Earlier this year, he also received his Chartered Financial Analyst (CFA) certification.

Speaking next will be Andrew Davidson, president of Andrew Davidson & Company, Inc., an investment management consulting firm that focuses on the application of technology to investment management. He was formerly managing director of mortgage-backed securities (MBS) research for Merrill Lynch and, prior to that, served at Exxon in its treasurer's department. His education includes an MBA from the University of Chicago in finance and a BA from Harvard College in mathematics and physics.

Mike Siegel is a vice president in the insurance industry resource group at Goldman Sachs. Mike says he does everything at Goldman that Chris does at Merrill.

* Mr. Anderson, not a member of the Society, is Director, Fixed-Income Research of Merrill Lynch & Company in New York, New York.

† Mr. Davidson, not a member of the Society, is President of Andrew Davidson & Company in New York, New York.

‡ Mr. Siegel, not a member of the Society, is Vice President, Fixed-Income Division of Goldman Sachs & Company in New York, New York.

Mike has both a master's and a doctorate degree in economics from the University of Michigan.

After my brief introduction, Chris Anderson will begin his presentation by describing the regulatory need and how the process has been approached. Andy Davidson will then explain the chosen measure that is the subject of this panel discussion and discuss the types of CMO variability that it attempts to capture. I'll then describe the scenario selection process that's being used in this model. Mike Siegel will then talk about the prepayment assumptions and the method to be used to collect prepayment data. Then Andy will talk about the formula and show some sample calculations. I will then discuss the relationship to cash-flow testing and identify some other issues related to the index. Finally, Chris will wrap up and describe the regulators' reactions to these results.

Why is this topic on the agenda? CMOs have been a fact of life for some time, and if you sat through a panel on prepayment assumptions that Randy Boushek moderated, you heard about some of the difficulties in modeling prepayment risk. Prepayment risk is clearly one of the sources of variability that we need to deal with. It is probably also the dominant risk in the CMO market, because CMOs largely exist to redistribute cash flows to different classes of investors by disproportionately allocating prepayment risk.

Refinancing has undergone such a transformation in its efficiency that there may soon be no uncertainty in projecting mortgage prepayments if interest rates go down. Mortgages will be prepaid. And if rates go up, they won't be. Mortgages are becoming very efficient in exercising their option. So prepayment risk is certainly one concern.

Innovation is another. Wall Street dealers are not only allocating prepayment risk to different classes in disproportionate forms, but are now also bundling other risks as well. An example is the coupon risk in the case of floater and inverse-floater tranches.

But the real reason for this topic is camouflage. That is, you can't tell by inspection (at least in any format that regulators tend to see), what a CMO holding is. A planned amortization class (PAC) tranche might be called Federal National Mortgage Association (FNMA) 1990-52 Class C. A floater might be called FNMA 1990-52 Class F, and an inverse floater might be FNMA 1990-52 Class V. Unless you are able to intuitively map the alphabet into different classes, you're stuck when assessing what types of risks are included in a portfolio of CMOs.

The role of this committee, then, is to provide some early detective work and to develop a tool for regulators to help discriminate some of these risks.

MR. CHRISTOPHER T. ANDERSON: As Dave said, it is possible to look at a statement and not fully understand what is there. So, about a year ago, the NAIC put together a CMO agenda. There are four points on the agenda, and I would like to describe the progress that has been made on each point: (1) disclosure in the 1993 statement, (2) disclosure of class types, (3) improved accounting, and (4) cash-flow variability.

First is disclosure on the 1993 statutory blank. Schedule D, part one of the statement will be broken up into four sections. A security will be listed in one of these sections: (1) One to four family residential securities, (2) structured loan-backed bonds (loan-backed bonds on which principal and/or interest is paid or passed through pro rata), (3) loan-backed securities (securities with other securities behind them), and (4) other bonds. This is the new structure for the 1993 statement that responds to the regulators' first question, "How can I find CMOs?"

Second, regulators wanted to know the class type of each CMO listed. Class types were defined by a joint venture between the Federal Home Loan Mortgage Corporation (FHLMC), the FNMA, and Merrill Lynch and include 14 different principal payment methods and 12 different interest-payment methods in addition to several other types. When you link them all together, there are 227 different combinations, because some are multiple combinations of different types.

How is this information made available to regulators? It's reported each quarter directly to the NAIC. Even though it is not on the statutory blank, regulators now have access to CMO class types for a large number of CMOs.

Next, regulators thought that there were shortcomings in accounting for CMOs. They were not comfortable that statutory accounting was as good as GAAP. Even beyond that, there was a sense among some that it could be even better than GAAP accounting. So, there was a study done by the CMO accounting team, headed by Walt Chossek from Northwestern Mutual.

About a half dozen insurance industry accountants served on the task force, along with an equal number of big six accounting firms. They submitted their report in June 1993.

The NAIC Blanks Task Force is meeting, and some initiatives are being proposed there to make CMO accounting more precise. First, the task force is recommending the adoption of either the prospective or retrospective methods; second, it proposes to disallow the use of the composite method; third, it is requiring prepayment projections to be made at least annually; fourth it is recognizing negative yield adjustments as capital losses; and last, it is including these losses in the interest maintenance reserve (IMR). We are not here to discuss this in any detail, but to make the point that these are steps that regulators are taking to improve the quality of accounting for CMOs.

The fourth point on the CMO agenda contains two elements. One is screening; the other is cash-flow analysis. By screening, I mean that regulators wanted a way to scan portfolios, to see which ones contain CMOs with greater-than-average cash-flow variability, in various interest rate scenarios. So, the objective of the flow-uncertainty index (FLUX) method is to assist regulators in screening portfolios, and to indicate when to focus more closely on cash-flow testing results.

The second point is that we need to focus on cash flows, as the other speakers will make very clear. We did not attempt to develop the ultimate test of cash-flow measurement or portfolio measurement. Instead, we believe we have developed a cash-flow-based screening tool for regulators.

The FLUX measure responds to the regulators' request that we provide a single, numeric value for each CMO class to express the variability of its cash flows under different interest rate scenarios.

The FLUX measure that we have devised is what we call an "open architecture measure." Basically, this means that we will publish the recipe for FLUX. The other speakers will discuss the FLUX ingredients in more detail. The recipe will be updated and published once a year so that anyone will be able to perform these computations. Anyone who can generate CMO cash flows, given a vector of prepayments and interest rates, can compute a FLUX score.

Finally, it is important to stress that FLUX is intended to be a relatively easy, straightforward, simple computation, not requiring complex or elaborate computations.

MR. ANDREW S. DAVIDSON: The measure that we chose to go with, after extensive testing, was something that I developed a calculation for, which is the "flow uncertainty index" or FLUX. FLUX is a score that we're going to assign to each of the individual bonds.

What does this measure? To justify coming up with something new, you have to first identify what's wrong with all the other measures that we have. If you look at a Bloomberg screen to analyze a bond, there's a full page describing all the different things that you'd presumably ever want to consider. For other analytical systems, there are also hundreds of different numbers that one can analyze. Why do we need another one? First, let's consider our specific objectives, then review the different types of measures that were available, and finally say how FLUX differs from those tools.

At one of our early meetings of this committee, we tried to write down a number of risks that we wanted to consider. Following is a short list of things we came up with.

1. Duration. That seems important. When you want to talk about how risky a bond is, duration is often going to be the most important measure. Duration measures the price sensitivity or the length of life of the cash flows of the security, and you would think that should be the primary goal. But, after some discussion, we realized that this was not correct. Insurance companies often have very long dated liabilities, and sometimes the most appropriate asset (the asset that's going to reduce risk the most) for that institution is an asset with a very long duration. We don't want it to appear that if you have long-dated liabilities and long-dated assets, you have risk. So, we decided that duration should come off of our list.

2. Credit. That's another really big risk. However, the insurance industry is already dealing with credit risks through other measures, and besides most CMOs don't have any material credit risk. They're mostly either triple-A securities or are guaranteed by the Government National Mortgage Association (GNMA), the FNMA or the FHLMC. So, clearly, credit wasn't a relevant risk.

3. Convexity. Convexity is a measure of the option features in a bond. How do you measure convexity? The way to measure convexity is to have a large option-adjusted spread model, where you generate thousands of interest rate paths and tie them to a prepayment model. But we wanted to do this consistently and simply across all insurance institutions. Unfortunately, different models give different measures for convexity. Although a bond with a high amount of convexity may be a risky bond, it's not really what we're trying to measure here.

4. Volatility. This is basically just another measure of convexity.

5. Spread to treasuries. How much does the market price of an instrument change as spreads change? That's a good measure for a trading portfolio, but what we're really looking at asset/liability managers. They have liabilities and they're purchasing assets. They're using the cash flows of those assets to fund those liabilities. What does spread to treasuries have to do with that? If you've bought your assets and you have your liabilities, the spread may change on those assets, but you're still getting the cash flows you need to meet your liabilities. So that wasn't really the crucial risk. It might be nice to look at those measures, because they are probably a good indication that something's changed about the market: there may be some problem with that asset, or it may be less liquid than you thought it was before. Once again, though, this measure doesn't get to the core of what we're trying to deal with. So we said, "let's get rid of these." We've now crossed out virtually every risk measure that was readily available.

We finally decided that we needed a measure that would deal with prepayment uncertainty and income timing. In different interest rate environments or in different prepayment environments we wanted to know how much the cash flows of the instrument will change. What is the chance that those cash flows will be sufficient to meet the liabilities, and how is that going to affect the timing of income? Suppose you were expecting some cash this year, and instead of getting it this year, you get it next year. If you have to borrow in the interim, there's a cost associated with that. Or, if you get the money now but you really don't need it for two or three more years, you may have a reinvestment risk. Those are the types of risks that are going to show up in a cash-flow-testing analysis. We wanted to deal with those.

We also considered several other approaches and although they all had merit, we decided that they didn't fit within the specific structure of what we were asked for. One approach is a value measure. Is the bond good or bad? Almost any bond has good value if you buy it at a low enough price, no matter how risky it is. If you pay nothing for it, you're happy for whatever cash flows you get. But how do we measure value meaningfully? For example, interest-only securities often have option-adjusted spreads in excess of a thousand basis points. Yet these are the same instruments that companies have lost the most money on during the past year. So just because the bond was cheap according to some model, it wasn't necessarily the best bond to have in a portfolio.

Finally, we debated between using the market value versus the book value of a bond. In the end, we decided that the measure should be independent of both the market value and the book value, because it is purely a measure of the cash flows and the stability or variability of those cash flows.

I have outlined the basic steps to calculate FLUX for a bond.

FLUX Process

1. Calculate bond cash flows for base case and five scenarios.
2. Use month-by-month prepayment speeds for bond collateral.
3. Calculate PV measure and timing measure for each scenario.
4. Add PV measure and timing measure to produce scenario flux.
5. Calculate root mean square of scenario FLUXes to produce bond FLUX.

We begin by calculating the cash flows of each individual bond in a base case and five additional scenarios. We're going to be using month-by-month prepayment speeds (for the collateral) to calculate the cash flows of the bond. Then we develop two different measures, a present-value (PV) measure and a timing measure, for each bond and each scenario. We add those two scores for each scenario to produce a scenario FLUX. Then, we combine the scenario FLUX scores by using a root mean square approach. This means that you square each number, take the average of the squares, and then take the square root. A root mean square allows you to capture both the average and the standard deviation of the variability in one calculation. If four numbers have the same mean, but they have a large standard deviation, they will have a higher root mean square than if all four numbers were exactly the same. So bonds that have higher scores in one scenario than another scenario will get a higher score than if they had been constantly "risky" across the scenarios.

I want to briefly describe the two major components of the analysis, the present-value measure and the timing measure. For the present-value measure, we calculate the present value of the cash flows of the bond, first for the base scenario, and then for each of the interest rate scenarios, all by using the same discount rate. Usually when you're doing mortgage securities analysis, you calculate effective duration or option-adjusted spreads by varying the discount rate for each scenario. But in this case, we hold the discount rate constant, because all we want to do is compare the total amount of cash flow. We could use a zero discount rate, at which the present value would be the total cash flow, but we thought that taking present values was consistent with the notion that the stability of early cash flows is more important than the stability of later cash flows.

Next, we compare each scenario present value with the base-case value, and we only consider negative changes. So if the bond has more cash flows in a scenario than in the base case, it gets a zero for that situation. If the present value of the cash flows is lower in a scenario than in the base case, then we take the difference and carry that through the analysis.

What this measure captures is the prepayment risk of premiums and the extension risk of discounts. The bonds that get higher scores because of this component are interest-only (IO), principal-only (PO), and other very high-premium or deep-discount bonds.

The second component, the timing measure, is a little more complicated. Suppose you have two bonds that both have a coupon equal to the discount rate, so the present value of the bond never changes under different scenarios. For example, if we're using a 6% discount rate, and a bond has a 6% coupon, under the present-value analysis, regardless of whether the bond is completely paid off tomorrow or ten years from now, it will still have the same present value of 100. But there's still a risk associated with that bond. The risk is that if the bond pays off tomorrow, and rates are lower and you have to reinvest, you have that reinvestment risk for the next ten years. Alternatively, if the bond extends to ten years, and you have to borrow because you need the money tomorrow, then you have to borrow at whatever interest rates are prevailing.

So the timing measure is used to measure reinvestment risk or asset/liability management risk. The types of bonds that will receive high scores on this component are support bonds and other bonds that have a great degree of average-life variability. Why not just use average-life variability? The problem with average-life variability is that it's just an average and fails to discriminate between a cash-flow barbell versus cash flows centered at one point. For example, consider a bond that is really a combination of two bonds, a one-year bond and a ten-year bond, with an average life of five years. If in one scenario it's a one and a ten year, and in another scenario it's a five-year bullet, you'd say the bond has zero average-life variability. But it still has a lot of cash-flow variability. You still have a very changing structure of cash flow. This timing measure is designed to capture that.

To do that, we use a technique called "cumulative present value," whereby we take differences in cumulative present values. Later I'll discuss those calculations, but in the meantime, let the other panelists describe the process of determining the scenarios and the prepayments, so you can first understand the inputs to the model.

MR. HALL: The work of the scenario selection committee was to define scenarios for the FLUX measure. This was a task that we undertook fairly late in the process, because we had to develop fairly clearly the type of measure we wanted before we could model the scenario process. So the ordering of this presentation, while it corresponds fairly closely to the way the project was mapped out on day one, isn't at all consistent with the sequence of events that we followed while actually developing our product.

Andy mentioned that FLUX was designed to measure five scenarios versus a base case. We didn't know that when we began building the measure, and in fact, the work of the scenario selection group was to identify how many scenarios were needed and what those scenarios should be. A primary goal was to minimize the number of scenarios to model. That was important to speed computation, because we expected that even when running this on fast computers, if you calculated this measure for many securities at one time, you'd burn a lot of CPU time. More importantly, limiting the number of scenarios would simplify the collection of prepayment rates, as Mike will discuss in a moment.

We wanted to use scenarios that were realistic, but also somewhat stressful. We didn't want to make them so stressful as to invalidate the results. Effectively, if you model prepayments too fast or too slow, all mortgage securities start to look the same. We hoped to find that middle ground, in which there was some meaningful discrimination of variability.

We thought it was important to include whipsaws, not necessarily because we thought whipsaws were necessary, but because regulators thought they were important to model. Because regulators were our intended audience, if we didn't include whipsaws, we were going to have to deliver a compelling counterargument.

We did not require mirror symmetry; that is, for every rising scenario we didn't have to have an equal but opposite falling scenario and vice versa. We also didn't need to use scenarios that were consistent with Regulation 126. In fact, for other reasons, we did not want to use those scenarios, because we didn't want FLUX to relate directly to cash-flow testing. Ultimately, rather than the actual choice of scenarios, the real objective of our committee was to develop a scenario selection process that could be used periodically to redefine new scenarios. One of the outputs of our group was to define a set of actual scenarios to be used in the first testing of FLUX, but the selection process is the true output of our committee.

To define a means for testing scenarios, we relied heavily on Tom Ho, who is the President of Global Advanced Technologies and also an acclaimed academician and a compulsive researcher at heart.

Tom approached this process by defining a spanning set of scenarios. He used a trinominal lattice framework, which means that at each time (or node) in your interest rate model that you assume interest rates can change, they can change only in one of three ways: up, down, or unchanged. To convert this to a workable number of scenarios, we assumed that interest rates could change in one, three, five, and seven years from today with all scenarios starting from current interest rates.

We assumed that movements of interest rates occurred linearly by 150 basis points per interval. To satisfy our realistic but stressful scenario constraint, an ultimate bound of no more than plus or minus 300 basis points from our starting level was imposed. For example, we did not test scenarios where interest rates ever rose by more than 300 basis points above the starting level, nor where they ever fell by more than 300 basis points. We thought that as you moved outside those bounds, there wasn't meaningful information to be captured.

For the actual testing, we chose a set of 17 scenarios. These included rising, falling, and whipsaw patterns. Our baseline scenario assumed no change in interest rates.

We modeled only parallel yield-curve shifts, those in which the long end and the short end are moving in a parallel manner at all times. That sacrificed some information. For example, a number of prepayment models include a yield-curve function. They depend on the level of interest rates at both the long and short end of the yield curve. Also, floaters and inverse floaters, by virtue of their coupon formula, can be influenced by the yield curve. However, we again thought that this was getting beyond what we needed to capture. In effect, most floating-rate securities were going to score low, no matter what. Inverse floaters were going to score high, no matter what.

And because we weren't trying to create a very subtle discriminating tool, parallel-yield curve shifts would satisfy simplicity but yet capture the dominant features that we needed to model.

To have a sufficiently large database for testing, we selected the first five agency CMO deals that were issued in each month during the period of January 1991 through December 1992. Interestingly, as of that time (the analysis was done in June 1993), those two years of issuance reflected 60% of all outstanding CMOs in the market. Our database included 2,039 individual CMO classes.

The process that we followed then was to calculate FLUX scores for each of the CMO classes. For each of the 16 (nonbase) scenarios, we calculated a FLUX score. We next identified the scenario for each bond that generated the maximum FLUX score. The intent was to detect those scenarios that caused the most stress for the greatest number of bonds.

Fortunately, five scenarios stood out, which seemed like a workable number. These are depicted in Chart 1. The percentages on the chart indicate how many of the securities tested generated their maximum FLUX score on that scenario. The scenario that caused the highest FLUX score most frequently was the down-300-basis-point-and-hold scenario. (Note that we're not declining the full 300 basis points until year three in our modeling.) It captured the maximum FLUX score for a little more than 37% of all the securities modeled.

CHART 1
Results: Five of 16 Scenarios Produced
97.3% of Maximum FLUX Scores

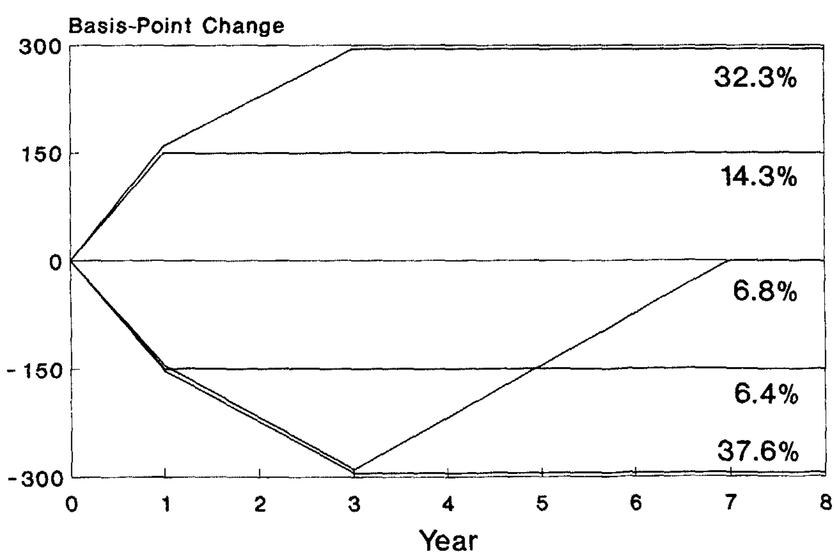

The next most flagging scenario was the one that went up 300 basis points, and that totaled 32%. These two combined accounted for roughly 70% of the maximum scores.

Interestingly, the next most flagging scenario was the one that rose 150 basis points during the first year and then leveled out. It managed to capture 14.3%. It's interesting that this scenario generated a higher score for some bonds than the one that allowed rates to keep rising.

The next two most common scenarios were the down 150 and a whipsaw (lo and behold, we did get a whipsaw scenario out of this, which was a happy result for the reasons I articulated), in which interest rates dropped 150 basis points during the first year and then another 150 during the next two years. During the ensuing four years, they grade back to the starting point. That whipsaw scenario captured almost 7% of the maximum scores.

Now, those five scenarios total 97.3% of the maximum FLUX scores out of all 16 scenarios modeled. The next closest scenario captured only 0.8%. In fact, there were seven scenarios that accounted for the other 2.7% of scenarios that we tested. Four of the scenarios generated no maximum scores. So, choosing the top five seemed like a fairly clean dividing line to draw. Any scenario that generated 6% or more of the maximum results was to be one of our model scenarios, and any that fell below the 1% level were out of scope. If results in future-years analysis sift out that way, that will be great. If we get more graded results, then someone will have to decide where to draw the dividing line. But it was very convenient that, in our analysis this first time around, we found a very clear breakpoint.

Having defined these five scenarios relative to our baseline scenario, we now have to model prepayment rates. And that's where Mike Siegel takes over.

MR. MIKE H. SIEGEL: I have the responsibility for the easiest part of this project, and that's developing the prepayment speed projections that are going to be used in the FLUX calculations.

To briefly recap, the ultimate goal is to develop a consistent measure of cash-flow variability for use by insurers and regulators. This requires three elements to ensure consistency. Those three elements are (1) the formula, which Andy Davidson will work you through shortly, (2) the set of interest rate scenarios that Dave Hall just discussed, and (3) the set of prepayment speed projections, which is my area of responsibility in this project.

The charge to the prepayment speed subcommittee was to develop a consensus estimate of mortgage prepayments for each scenario, for each collateral type. The question is, how do we develop this consensus forecast of prepayment projections per collateral type per interest rate scenario? We considered a number of alternatives.

The first thing we did was consider ignoring the request. And we're still thinking about that. That, however, does not satisfy the regulators. The second thing we considered was to make the projections up, and I think that is still under consider-ation. Then, quite seriously, we thought about using just two numbers: something like 50% prepayment speed assumption and 1,000% prepayment speed assumptions. Fifty percent prepayment speed assumption represents a slow speed, and that would capture all of the extension characteristics for the bonds, and a 1,000% prepayment speed assumption would be a very fast speed, which would show the shortening sensitivity of the bonds.

We had two other alternatives; we could develop a model ourselves or we could get help. We chose help.

We have asked for the assistance and support of the Public Securities Association (PSA). For the last several months, we have been working with a research subcommittee of the PSA, in order to have member firms of the PSA develop and contribute these prepayment speed forecasts. In essence, it's not any individual firm's responsibility, it's everybody's collective responsibility.

We have been getting a lot of support from the PSA. It is in the process of working with us to develop final guidelines, which I will discuss very briefly. The research committee will meet at the end of the month to develop its recommendations for the full operating committee of the PSA, and we hope to get its final approval by November 1993.

We hope the PSA will ask its 300 member firms to prepare monthly prepayment speed forecasts. This will be done probably once a year, and it will be done as of a specified date. We've been talking about using September 30, which is the end of the 3rd quarter. For this year, we've missed that date and we'll probably end up using a date like December 1. But on a going-forward basis, at every September 30, the individual firms will take the prevailing yield curve and the projected interest rate scenarios and use them to forecast prepayment speeds by collateral type.

These prepayment-speed forecasts will be done on a monthly basis. They will go out 360 months (or as far out as the pool of collateral goes). Rather than one single speed, we'll receive 360 monthly speeds if the pool goes out that far.

The collateral types will be those that are on the PSA telerate screen. The PSA Research Committee reviews that screen semiannually, and if new collateral types come into the market, they will then vote on having them added to the screen.

Currently on the screen are virtually all of the 30-year pass throughs, both GNMAs and conventional mortgages (FNMAs and FHLMCs together). They include new production, seasoned, and what's known as short weighted average maturity (WAM) pools. They also include 15-year mortgages and five- and seven-year balloon mortgages.

In addition, the members will be asked to provide prepayment speed forecasts for nonagency collateral. In other words, not FNMA, FHLMC, or GNMA qualifying collateral, but collateral that backs nonagency transactions.

Finally, the PSA will collect all these speed projections from its members for the scenarios and for each of the different collateral types, and we'll develop what's being called an NAIC speed. It is not an official NAIC speed, but it will be the speed used in the FLUX calculation. That will be the median speed per collateral type, per scenario, per month. In other words, for the month of January 1994, for GNMA 8s with a 330 WAM, we will take the median speed from all of the forecasts.

What happens for collateral that projections haven't been made for? For the most part, there will be a rule that will provide a mapping for that collateral back to similar collateral in which prepayment speed projections are being made. For example, if you have a quarter coupon, something like an FNMA 6, the rule will be to go to the next closest coupon, or to round up. So, if you have a 6.5% coupon, use a FNMA 7%, prepayment projection. If you have a WAM that falls in between two projected WAMs, use the closest, and if it's in the middle, use the speed for the higher WAM.

Some CMOs use mixed collateral. For example, they combine both FNMA 7s and FNMA 8s. The rule will be to use FNMA 7 projections for the FNMA 7 collateral, and use FNMA 8 projections for the FNMA 8 collateral. You should not use a FNMA 7.5 speed projection.

In summary, our recommendation to the NAIC committee will be to use the support and assistance of the PSA, to have the PSA members develop their prepayment speed forecasts for the different collateral types according to the scenarios, collect the information, and consolidate it. This will produce one set of projections that will then be used for the FLUX calculations.

MR. DAVIDSON: We've had a fair amount of discussion about the process and the scenarios, prepayment speeds and some of the underlying characteristics of what we were trying to accomplish. Now I'd like to show you these calculations.

Our goal was to create an open system, in which the actual calculations are open. The formulas may look complex, but they can easily be programmed on any spreadsheet. The hardest part will be to project the CMO cash flows, unless you have a CMO cash-flow generator. Assuming you have these cash flows projected, there are six equations that make up the FLUX calculation (see below). There are actually two different FLUX calculations, one for fixed-rate securities and one for floating-rate securities. I'm going to focus on the fixed-rate security calculation.

FLUX CALCULATION FORMULAS

- Cumulative Present Value

$$CPV_{s,m} = \frac{Principal_{s,m} + Interest_{s,m}}{(1+r)^m} + CPV_{s,m-1}$$

- % Decrease in Present Value

$$\Delta PV\% = \frac{MAX(0, CPV_{base,M} - CPV_{s,M})}{CPV_{base,M}}$$

- Absolute Scaled Differences

$$ASD_{s,m} = \frac{CPV_{s,m}}{CPV_{s,M}} - \frac{CPV_{base,m}}{CPV_{base,M}}$$

- The Timing Score

$$T\%_s = \sum_{m=1}^{M} ASD_{s,m} \times V$$

- Scenario Flux

$$FLUX_s = \Delta PV\%_s + T\%_s$$

- FLUX for a CMO

$$FLUX = \left(\frac{1}{S}\sum_{s=1}^{S} FLUX_s^2\right)^{1/2}$$

The key to this whole measure is the concept called "cumulative present value." It's not something that we deal with very frequently. To calculate the cumulative present value for scenario *s* and for the given month *m*, sum the principal plus interest cash flow for that scenario and that month. Then discount it back to today, at the given interest rate, and add it to last month's cumulative present value. It's the present value of the cash flow from month *m*, plus the cumulative present value for month $m-1$.

If you think about cumulative present value, the value for month *m* (where $m = 1 \ldots M$, and *M* is the last month), is the present value of the bond for a given interest rate scenario.

The second equation is the percentage decrease in the present value. Remember, I said we're only going to include negative changes. Thus, it is the maximum of the difference between zero and the base case cumulative present value (the present value of the bond) minus the present value of the bond in the scenario, divided by the present value of the bond in the base case. This division is to convert everything to a percentage basis. Because we're dealing with cash flows and not prices, we've got to normalize everything in some way. So we divide through by the base case present value.

To calculate our timing measure, we develop another new concept. This is absolute scaled differences. To calculate this, I take the cumulative present value for that scenario in that month, divide it by the total cumulative present value, subtract that same ratio for the base case, and then take the absolute value of that difference. This measures whether I'm getting cash flows sooner or later. Taking the absolute value means that I don't care if it's sooner or later. Either way, it is something that I have to be concerned about.

Graphically, consider Chart 2 and Table 1. The top graph of chart 2 shows the cumulative present values of the base-case scenario. Now look at scenario 1. We can see that scenario 1 has a zero difference in the first year. There's a slight difference of about 0.5 in the second year, and then a difference of zero after that. So an absolute scaled difference only appears in year 2. That's the only time when I've gotten some cash flow a little bit sooner than in the base scenario.

In scenario 2, the cash flows were delayed by a year. I have a zero difference in the first year, a zero difference in the second year, and a difference almost equal to 1 in the third year. But by the fourth year, I'm even again. I received all the cash flows from both bonds. So you can see that scenario 2 is going to have a larger timing cost than scenario one.

To calculate the timing score, I sum the monthly absolute scaled differences multiplied by a volatility factor (*V*), for which we're using 1.5%. The volatility factor we use is going to be similar to the step size in the scenario selection.

We then calculate a FLUX score for each scenario by summing the PV score and the timing score for that scenario. Then we have to combine the scenario scores by using the root mean square calculation. You take the average of the square of each scenario FLUX, then you take the square root. Table 1 shows a sample FLUX calculation.

CHART 2

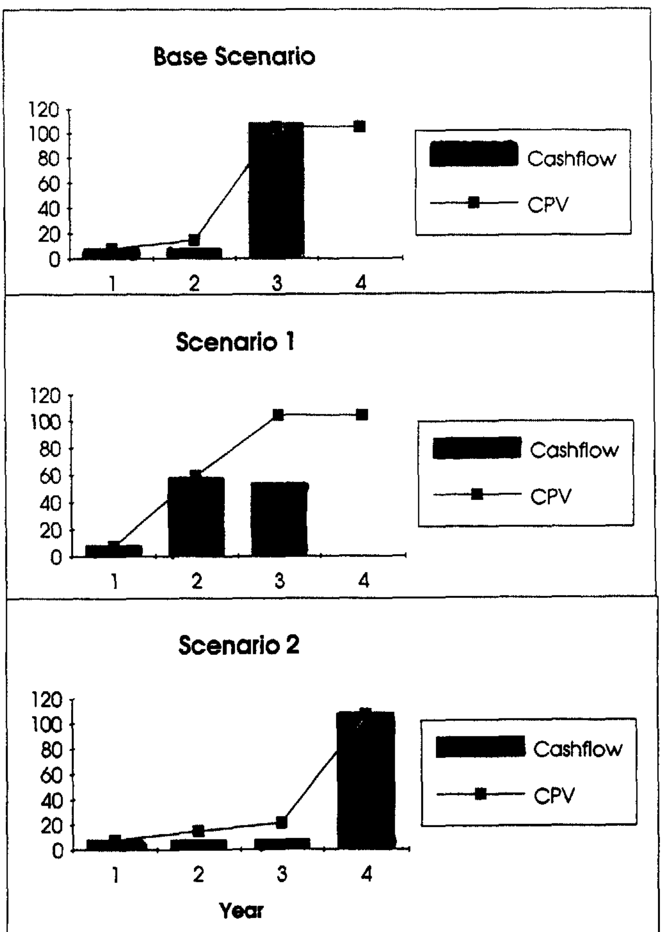

TABLE 1
FLUX Calculation Example
Interest Rate: 6% Volatility: 1.5% FLUX: 1.30%

BASE	1	2	3	4
Principal	0	0	100	0
Interest	8	8	8	0
CPV	7.547	14.667	105.346	105.346
Total CPV:	105.346			
Scenario 1	**1**	**2**	**3**	**4**
Principal	0	0.00	0	100
Interest	8	8	4	
CPV	7.547	59.167	104.506	104.506
ASD	0.001	0.427	0.0	0.0
Total CPV: 104.506 Change PV%: 0.80%		T%: 0.64% FLUX: 1.44%		
Scenario 2	**1**	**2**	**3**	**4**
Principal	0	0	0	100
Interest	8	8	8	8
CPV	7.547	14.667	21.384	106.930
ASD	0.001	0.002	0.800	0.00
Total CPV: Change PV%:	106.930 0%	T%: 1.20% FLUX: 1.20%		

The floating rate security calculation is very similar. The main difference is that there is no penalty for the cash-flow variability associated with the floating-rate coupon changing as expected. If you buy a London Interbank Offered Rate (LIBOR) floater and it floats with LIBOR, that's probably a good thing, and that type of variability should not be penalized. (On the other hand, if you've made a mistake and you bought a floater when you meant to buy a fixed-rate bond, you probably should be penalized. But the measure is assuming that if you're buying a floater, you know that you're buying a floater, and if you are buying a fixed-rate bond, you know you are buying a fixed-rate bond.)

Table 2 shows median FLUX scores from our sample universe by bond type. The bond type tells you how the bond was structured. It identifies the nature of the rules that distribute principal and interest from the collateral. The actual variability of the bond is going to depend on how that type fits into the overall structure and also on the variability of the cash flow from the collateral.

TABLE 2
Median FLUX Scores

Payment Types	fix	Z	inv	io
AD	1.6			16.5
PAC	1.8	6.5	8.9	27.6
TAC	2.0	7.4	14.0	26.2
SEQ	1.7	8.9	8.7	27.6
SUP	4.2	17.0	19.5	9.8

We show here only the median scores, but in many cases the actual variability within each category is extremely large. In virtually every category, we had some bonds with a FLUX score close to zero. If you have a bond that's on the verge of paying off, no matter what it was two or three years ago, it probably doesn't have much risk left. It may have had a lot of risk before, but going forward it has very little risk.

We have shown here five different principal payment types: accretion directed (AD), planned amortization class (PAC), targeted amortization class (TAC), sequential (SEQ), and support (SUP). We see that these bonds have relatively low scores if they have fixed-rate coupons. But we can also see that the support bond, the one that has the greater average life variability, gets a median FLUX score which is about three times as high as that of the PAC bonds or the AD bonds. That's not to say that in our sample we didn't have plenty of PAC bonds with higher scores than support bonds.

Z bonds (accrual bonds) tend to have greater risk than fixed-coupon-paying bonds (fix). Inverse floaters (inv) tend to be on the order of magnitude of the Z bonds, although we see a lot more variability within those scores. That will depend a lot on the specific structure of the inverse floater, and in particular, on the coupon leverage ratio for that inverse floater.

Finally, the IO securities tend to have the highest scores. The reason for this is that they have a tremendous potential to lose present value in a fast prepayment scenario. That's because the actual amount of cash flow can decline quite precipitously.

These are just median scores from a particular sample. But I hope they give you some flavor of the magnitude of the numbers and the relative scores for different types of bonds.

MR. HALL: Now that Andy has described what FLUX is, I'm going to describe what it is not. And I think the first thing to do is talk about how FLUX relates to cash-flow testing.

To sum it up, the direct relationship I see between the FLUX calculation and cash-flow testing is that there is no direct relationship. FLUX is designed solely to be a regulatory screening tool to detect relative cash-flow variability. For any actuary who can model CMO cash flow correctly for cash-flow testing, I have trouble seeing that the FLUX calculation can provide any useful information on a portfolio.

However, I think there are some considerations in cash-flow testing that relate to FLUX. Clearly, a portfolio of CMOs that all score on the low end of the FLUX scale may indicate that more simple CMO models are sufficient for cash-flow testing. I don't mean overly simple. I don't mean you can assume there's no variability, but a model that captures the general drift of cash flows may be sufficient. This is probably the most dangerous assertion I've made. And to rely on that may be somewhat misleading.

The safest statement I think I can make is that a high FLUX score indicates that more highly developed CMO modeling should be critical for your cash-flow testing. It indicates that in order to do this right, you need a model that is capable of calculating deal-specific cash flows that relate to the actual structure of each particular bond in your portfolio.

The biggest consideration is that a high proportion of high FLUX scores may trigger greater regulatory scrutiny of your cash-flow-testing results. I compare this to a tax audit. It's been well documented that if, on your tax return, you fall outside of certain reporting norms, for example, on your charitable deductions or your medical expenses, this increases the likelihood that your tax return may be audited. And such is the case with FLUX. If your portfolio exhibits a lot of high FLUX scores, that's going to increase the likelihood that a regulator will be interested in your cash-flow-modeling capabilities. That does not mean that the regulator finds a problem with your securities, but it does mean that the regulator will be concerned that you have the capability to appropriately model your cash flows.

When I first became involved in this project, one of my intents was to make sure that we calculated something that could be applied at the portfolio level. As any sophisticated mortgage-backed securities investor knows, many risks can be offset within a mortgage portfolio, such that the combination of two offsetting securities can provide a more stable profile than either security individually.

Although that was one of the biases that I had going into this process, I backed away from that stance for what I think are valid reasons. FLUX, quite frankly, is based on highly deterministic correlations in risk factors. For example, we only consider parallel yield-curve shifts. We also assume a very tight correlation between the level of yields in the scenarios and prepayment speeds. And we assume a high correlation between prepayments of dissimilar collateral; for example, prepayments on GNMA 8s versus FNMA 10s. If we develop a portfolio-based FLUX, we would have a significant risk of masking the reliance on these correlations. Clearly, we couldn't rely on only five scenarios to capture these correlation risks. We've been able to take shortcuts by developing only a bond-by-bond measure.

My concern is that if we look at a portfolio measure, where we first aggregate the cash flows of an entire portfolio and then calculate a FLUX score, we have a danger that we may calculate a low FLUX score, when there really is a great deal of risk that the prepayment correlations within the FLUX measurement will not occur. By virtue of not defining a portfolio application for FLUX, we hopefully avoid misleading results. If you have portfolios with offsetting risks, they really need to be modeled by using more sophisticated analytical tools. FLUX was not designed to capture that type of risk. And it really should not be used for that purpose.

FLUX is not correlated with what I'll call CMO morality. A high FLUX score does not indicate that you have a bad bond, nor does a low FLUX score indicate a good bond. We've used some terms fairly loosely, like "risky," in describing high FLUX scores. A high FLUX score does not necessarily indicate a risky bond. In some cases, the variability may be very beneficial, and FLUX, while attempting to discriminate between negative variability and positive variability, nonetheless cannot fully distinguish all of those issues. More importantly, high FLUX does not indicate that a bond is inappropriate for a portfolio. Nor does low FLUX indicate that a bond is appropriate. Most floating-rate bonds, for example, will generate a low FLUX score. If a company is using these to back structured settlement contracts with very long durations, that's an incredibly poor match, in spite of the fact that these CMOs don't have a lot of variability.

Finally, FLUX is totally independent of price or book value. Every CMO offers good value if the price is low enough. And similarly, every CMO has bad value if the price is high enough. We know that if we can get something for free, and the only uncertain thing that can happen is that we'll receive some cash flows at some time, (we just don't know how much and when), then free is a cheap price for that bond. We could argue about the best threshold price that discriminates good value from bad value, but there is clearly some price that offers good value. And similarly, every bond, no matter what its redeeming features, can be overpriced. Again, FLUX scores do not have any correlation to this type of valuation measure.

So, the bottom line is, we've created a single-purpose tool to be used only as a screening device by regulators, but not for anything else. That's probably one of the central points we want to leave you with. That is, you should not try to use this FLUX calculation for other purposes.

Having gone through this whole process, I'm now going to give Chris a couple of seconds to recap what the regulators think of this tool that we've developed.

MR. ANDERSON: We have kept regulators advised of our progress throughout the process to be sure they are fully aware of where we are. To the extent that we have not completely resolved all issues regarding prepayment speeds and other minor issues, regulators are certainly aware of the specifics of this as well.

In summary we can ask -- who can do this calculation? Anybody. It's an open architecture system. Your companies can do it. Third-party vendors can do it. What? They're computing a single numeric score. Where? On your own systems or on systems that you use in the investment area. You can program it yourself because it is a relatively straightforward calculation. When? Once a year. All of the inputs will be released by the Invested Asset Working Group and, as you learned, prepayment speeds will be developed as of September 30, 1994. We would suggest, however, that companies will also probably project prepayment speeds, interest rates, and volatility measures. How will they do it? That is up to you. We have created a method that is in the public domain. The last question is why? As professionals, I would ask that if you don't have at least one answer to that question by now, please don't claim continuing-education credit for this session. But I'll leave that up to you.