1998 VALUATION ACTUARY
SYMPOSIUM PROCEEDINGS

SESSION 3PD

ASSET MODELING CONCEPTS

Frederick W. Jackson, Moderator

Wesley Phoa

ASSET MODELING CONCEPTS

MR. FREDERICK W. JACKSON: Wesley Phoa is our panelist. He has a Ph.D. in mathematics and is director of research for Capital Management Sciences (CMS). He will talk first about techniques for modeling real estate asset-backed securities. After that, he'll discuss measuring and managing option risk.

DR. WESLEY PHOA: The first presentation I'm going to give is on modeling real estate asset-backed securities, and that includes home equity loans, home equity lines of credit, manufactured housing, and so on. The second half of my talk is on measuring and managing option risk. It will go along an entirely different track.

The main subject of the first half of the talk on real estate asset-backed securities, involves modeling prepayments on assets such as home equity loans and manufactured housing loans. As you're probably aware, asset-backed securities like this form a relatively new market. They've been actively issued and actively traded for only three or four years. The market in home equity loan-backed securities has really exploded in the past two years or so. As late as 1996, bonds like this were being sold to investors on the story that they had stable cash-flows, that they were AAA-rated bonds which essentially had very predictable prepayment characteristics. What we've seen since then is that, like any other kind of mortgage, a home equity loan doesn't really have particularly stable cash-flows. Prepayment risk, optionality, and the degree of optionality in a home equity loan, although it's a lot less than traditional conventional mortgage or a Federal Housing Administration (FHA) mortgage, are all important. They form an important part of the risk of a bond portfolio. In particular, something that wasn't so obvious back in 1995 or 1996, but that is very obvious now is that if you own a home equity loan-backed security or a manufactured housing loan-backed security, then you own something that does have negative convexity or less negative convexity than a conventional mortgage. It is still something that has to be taken into account.

There are certain problems, though, in actually trying to capture the kind of optionality embedded in assets like this. In particular, there's very little historical data available, and the market in these assets themselves is changing quite rapidly. What I wanted to do is go back to first principles and talk about loans like this and actually discuss what prepayment modeling really is. I think it has been traditional to regard prepayment modeling as some fairly complicated exercise in statistics or in econometrics, taking a huge mass of prepayment data, historical data, and fitting some more or less complicated model to that. In this particular case, it's much more useful to take a step back from that and to look at it on a more fundamental level to see what it involves. The right perspective to start thinking about prepayment modeling is to identify the major causes of prepayment and to quantify the impact of each of those causes. So prepayment modeling actually starts on quite an intuitive level. The more sophisticated statistical tools only come into the process later on.

The major causes of prepayments will depend on the type of loan that you're looking at. The causes of prepayments on a home equity loan made to a B or C borrower might be different, or at least different in emphasis, from the causes of prepayments most relevant to a pool of conventional mortgages. Just to give you one example, historically curtailments or partial prepayments have been a very, very minor part of prepayments on conventional mortgages, and there are various reasons for that. Tax reasons, for example, make it not a particularly efficient way of saving, whereas on home equity loans, particularly home equity loans with shorter maturities, curtailments have been a much more important source of prepayment. An important reason is that the interest rates have been much higher than on conventional mortgages.

After identifying the major causes of prepayments, you have to go on to quantify the impact of each cause. That's going to depend on the market environment. For example, if interest rates fall, refinancings become more important. That's where the optionality comes from The important thing here is that the whole analysis is not purely quantitative, which is unfortunate. It combines a fairly rigorous analysis of the data with a lot of market judgment, and that means judgment not just about capital markets but about the market environment in which borrowers live and the changes in lenders' practices in the competitive environment that borrowers are facing. So that qualitative element, which is a bit unsatisfactory, unfortunately seems to be unavoidable when you're doing prepayment modeling.

Let me just give a slightly more systematic list of the major causes of prepayments. They fall in two broad categories: prepayments that are not interest-rate-sensitive and prepayments that are. Non-interest-rate-sensitive prepayments give you a baseline cash-flow profile for a security, and interest-rate-sensitive prepayments give you a baseline cash-flow profile for a security, and interest-rate-sensitive prepayments give you a baseline cash-flow profile for a security, and interest-rate-sensitive prepayments are what give rise to the optionality and the negative convexity of the security. Non-interest-rate-sensitive prepayments are those that arise from ordinary housing turnover. For example, there might be a mortgage which is due on sale, which nearly all are, and refinancing, which is due to the changing credit of the borrower. Let's say you're a borrower with B or C credit. You've taken out a home equity loan at, say, 12–14%. As your credit improves, there's a big incentive to refinance your loan. Mortgage prepayments may also be triggered by a desire for the borrower to release equity in a property. For example, if you have built up $30,000 or $40,000 of equity in your house, you might want to release some of that equity to pay for your kids' college education. That creates an incentive to refinance the mortgage and raise the loan balance. Then there are curtailments, which I spoke about before, which are a form of saving; and less important factors, like default, death, and destruction of the house, make up the balance of observed prepayments.

FROM THE FLOOR: Could you define curtailment?

DR. PHOA: A curtailment occurs when a borrower makes a repayment that doesn't pay off the whole loan but makes a repayment that is higher than the monthly payment required. If the monthly payment is $1,000, and the borrower pays down $1,500, then the extra $500 is applied to pay down the principal, and that reduces the remaining term of the loan. That's why it's called a curtailment.

The other broad category of prepayments are interest-rate-sensitive prepayments, and we tend to identify those with refinancing, but there are other kinds of interest-rate-sensitive prepayments that occur. For example, what we've seen in the past 12 months or so is increased housing turnover when loans become more affordable; as such, borrowers can afford more expensive properties and can trade up. In the 1980s, we saw quite a bit of increased levels of defaults on adjustable rate mortgages when monthly payments went up dramatically as interest rates rose.

Once we've pinned down the different sources of prepayment, we can try to quantify each cause of prepayment in turn. One advantage of distinguishing between different sources of prepayments is that it gives you a lot more flexibility in using different sources of data to quantify different causes of prepayments. For example, if you're looking at a specific mortgage pool, you might expect that specific factors might affect how borrowers refinance out of that pool. The level of housing turnover is going to be driven by housing turnover in the economy as a whole.

In any case, the traditional way of quantifying prepayments is to look at the historical prepayments and to conduct some historical or cross-sectional analysis of historical data, preferably data that are available at the loan level, although it's very uncommon to be able to get detailed data of that kind, and to fit that to some kind of model. It's important to remember, though, that historical prepayment data aren't the only kind of data that are relevant to fitting a prepayment model or to estimating a prepayment model. Other kinds of statistical data can be very important. For example, there are historical housing turnover data, information on consumer indebtedness, and information from all different kinds of sources. In particular, it's not just quantitative data, and it's not just time series data that are relevant when you're trying to estimate a model For example, let's suppose you know for a fact that the competitive environment is changing for a particular kind of loan. Let's suppose you know, as we do, that in the past 12 or 24 months, there have been a lot more participants in the home equity loan market, more competition for borrowers than there were in the earlier days of home equity loans. Then you'd have to use that information somehow to weight the relevance of recent prepayment data versus older prepayment data. That's an example in which exercising some judgment that isn't purely quantitative is very important.

Now let me give a different example of where qualitative information comes into the process, and this has to do with estimating prepayments on particular kinds of home equity loans, those aimed at lower credit borrowers' season. Just to remind you, seasoning of prepayments refers to the fact that if you start off with a mortgage pool, other things being equal, prepayments tend to be slow shortly after origination. Then, as the pool ages or seasons, prepayments tend to accelerate, and after a certain time they level off. With a pool of conventional mortgages, roughly speaking, prepayments start off slow, accelerate for the first two or three years—traditionally it's 30 months—and then they level off. What we see with home equity loans aimed at B and C borrowers is that seasoning appears to ramp-up over a much shorter period (about 12 months). However, the data are actually very noisy, and that's not really a statistically robust conclusion. The question is, should you build that observation into a prepayment model? The answer is, you should if you can think of a reasonable explanation for it because whenever you see something in the data, you should realize that it is occurring for a reason. If it is, it's just an artifact. If you see something, you should try to think of a reason. In this case there is a reason why there is an approximately 12-month ramping period for prepayments like this. That's because the main cause of prepayments on these kinds of assets is credit-driven. It's the fact that borrowers are initially getting these home equity loans at quite high interest rates. Their credit is improving over time as they're meeting their monthly payments regularly. Once you've reestablished a good credit history and 12 months clean payment history, that is usually enough. That puts you in a very good position to refinance either into a cheaper home equity loan or even into a completely new mortgage.

Let me just give you an example of prepayment seasoning. I've taken some prepayment data on UCFC home equity loans (see Chart 1). They've been aimed most often at lower credit borrowers. They have fairly high fixed interest rates. There is a clear pattern of seasoning over time—prepayments accelerating over the first 12 months or so. What's happening here is that you get a fairly steep acceleration for 12 months—that's a function of credit-driven refinancings. You still get some acceleration even in the second year. That's probably more related to housing turnover. The data are quite noisy.

I wanted to show you a couple of other graphs to do with interest-rate-sensitivity (Chart 2). This is trying to measure something else. I've estimated the other component of a prepayment model which is interest-rate-sensitive prepayments, and I've also left out some details. The basic method was to

CHART 1
Analysis of Seasoning: UCFC

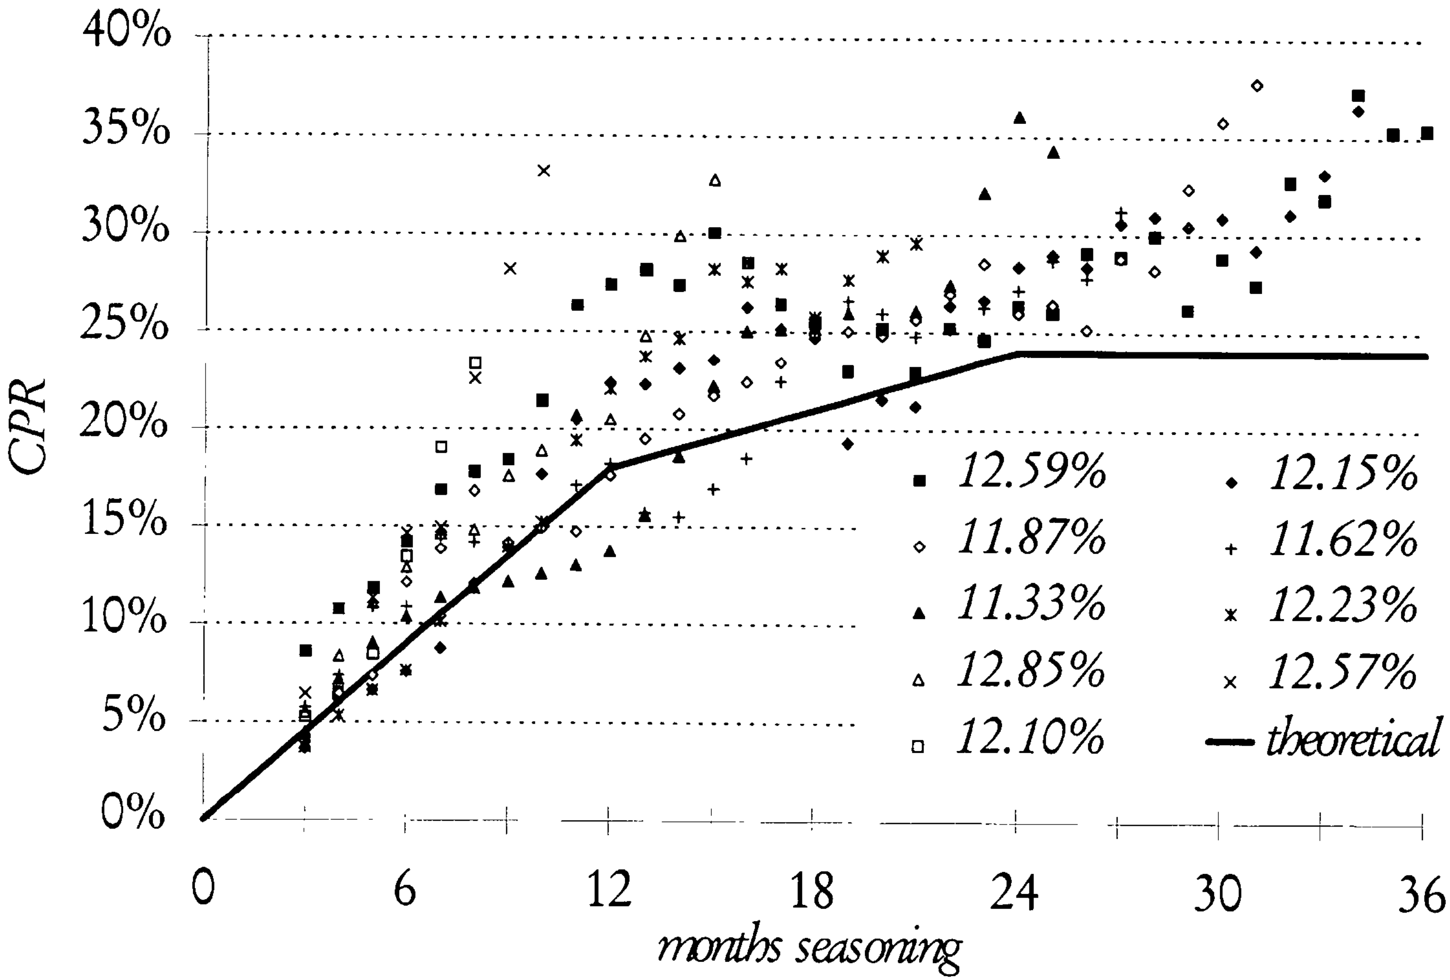

first estimate a non-interest-rate-sensitive, baseline, prepayment profile and then factor that out. You would find what is left and plot that against the interest rate differential between the coupon on the pool and the current Treasury yield. What we can see are two quite noticeable things. First, there does seem to be a fairly clear relationship between interest rates and prepayments. When interest rates are lower, in other words, when a spread between the coupon and the current Treasury yield is higher, then prepayments are quicker. That's what you would expect to see, and that's what you do see. There's optionality in prepayments on home equity loans. When interest rates fall, prepayments rise.

The other thing that's extremely important to note here is that the data are very noisy. There seems to be a lot of variation between different pools, and even within a particular pool. Although you can see a relationship, you wouldn't place too much confidence in any of those particular regressions. What that tells you is first that optionality exists, and it's important to model it, but that you shouldn't place too much trust in any specific model that you happen to build.

Let me just give a very quick outline of the home equity prepayment model that we built at CMS. It's the model that's operating in our BondEdge product at that moment. We tried to keep the structure as simple as possible while capturing the different major categories of prepayments. We've modeled non-interest-rate-sensitive prepayments by breaking them into three pretty broad subcategories (Chart 3): prepayments due to relocations and credit-driven refinancings, which have a certain seasoning profile over time, which seem to be fairly constant over time, and then defaults that have a kind of humped profile which is what you would see empirically. As for interest-rate-sensitive prepayments, there's also a separate response function that just looks at an interest rate differential and estimates the kind of additional CPR due to refinancing and increased turnover and so on based on that (Chart 4).

CHART 2
Interest-Rate-Sensitivity

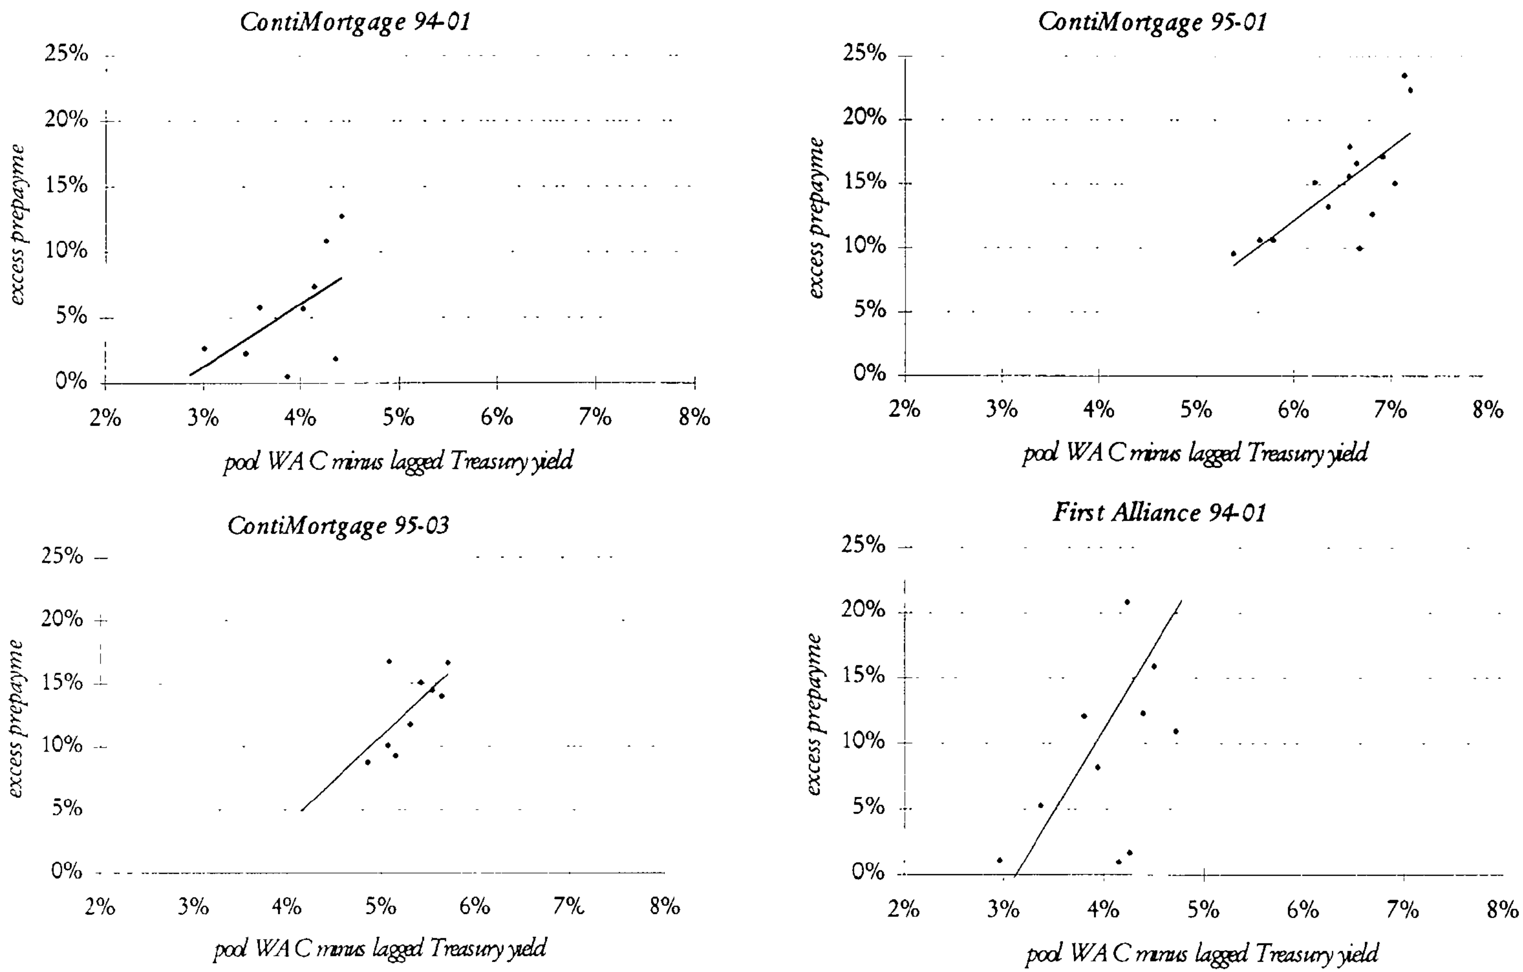

CHART 3
HEL Model: Non-Interest-Rate-Sensitive

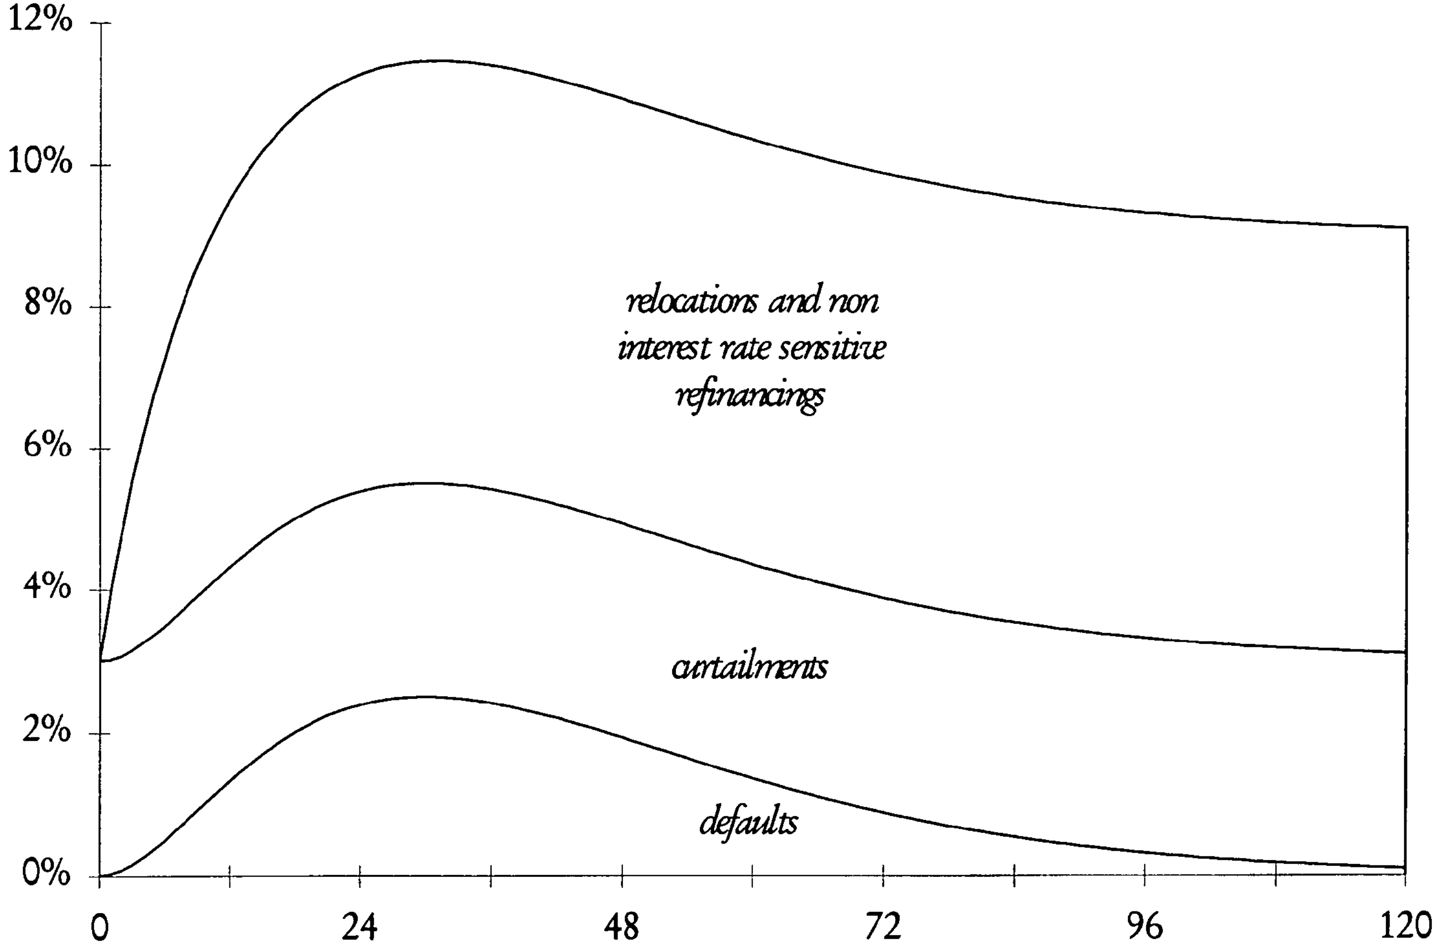

CHART 4
HEL Model: Interest-Rate-Sensitive

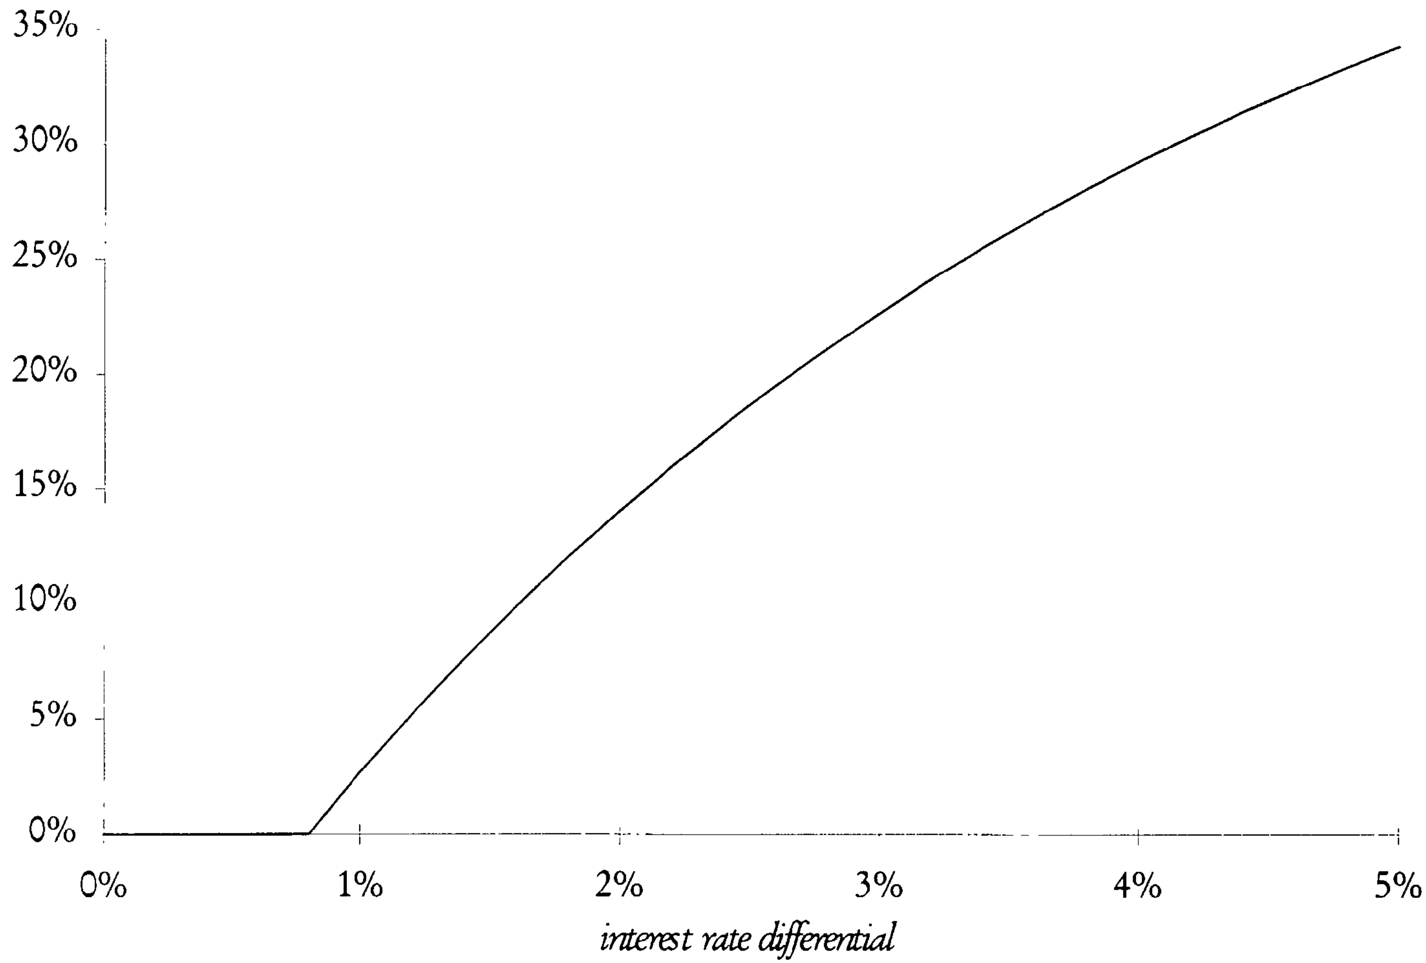

FROM THE FLOOR: What does the y-axis represent?

DR. PHOA: The y scale is an additional percentage conditional prepayment rate (CPR), over and above baseline prepayments.

For example, if we look at this graph and if the interest rate differential is 3%, then we're assuming that you're getting 20% of the mortgages refinanced or prepaid at an additional 20% prepayment.

We try to look at how well it fits the data. It's not very surprising that the model fits the data well because we constructed the model using those data in the first place. It's at least reassuring. It's not always so obvious from broker research that that's the reason you obtained such a good fit. One thing that Chart 5 does tell you, though, is that although it's an in-sample fit, at least the structure of the model itself was rich enough to capture what you saw in the data, and that's not a trivial observation.

CHART 5
Model Fit: Historical
UCFC 93-D1 Prepayments vs. Seasoning

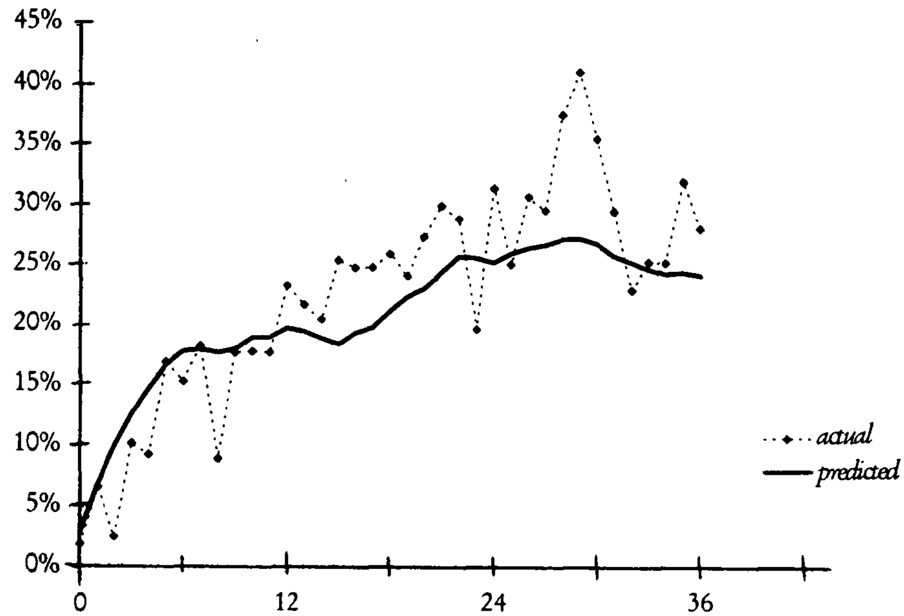

If you take just the time series of prepayments on one specific pool, which is a fairly early one, you're getting quite a good fit to the data. What's just as important is you get a reasonably good fit if you look at the data cross-sectionally rather than in terms of a time series (Chart 6). It's less common to see graphs like this, but here I've taken a series of observations on different home equity lines of credit pools, but at one particular period of time. You look at a single date, but you look at pools with varying degrees of seasoning and varying coupons. You can see that there's a pretty good correspondence between the observed prepayments and those that would have been predicted by the model as well. It's important to look at these cross-sectional fits, too. I should also say that our model has actually started to drift away a little in 1998. We're in the process of refitting it now, and, of course, it's always necessary to continually be examining how reliable your model is and refitting it.

CHART 6
Model Fit: Cross-Sectional
Equicredit Prepayments 1/31/97

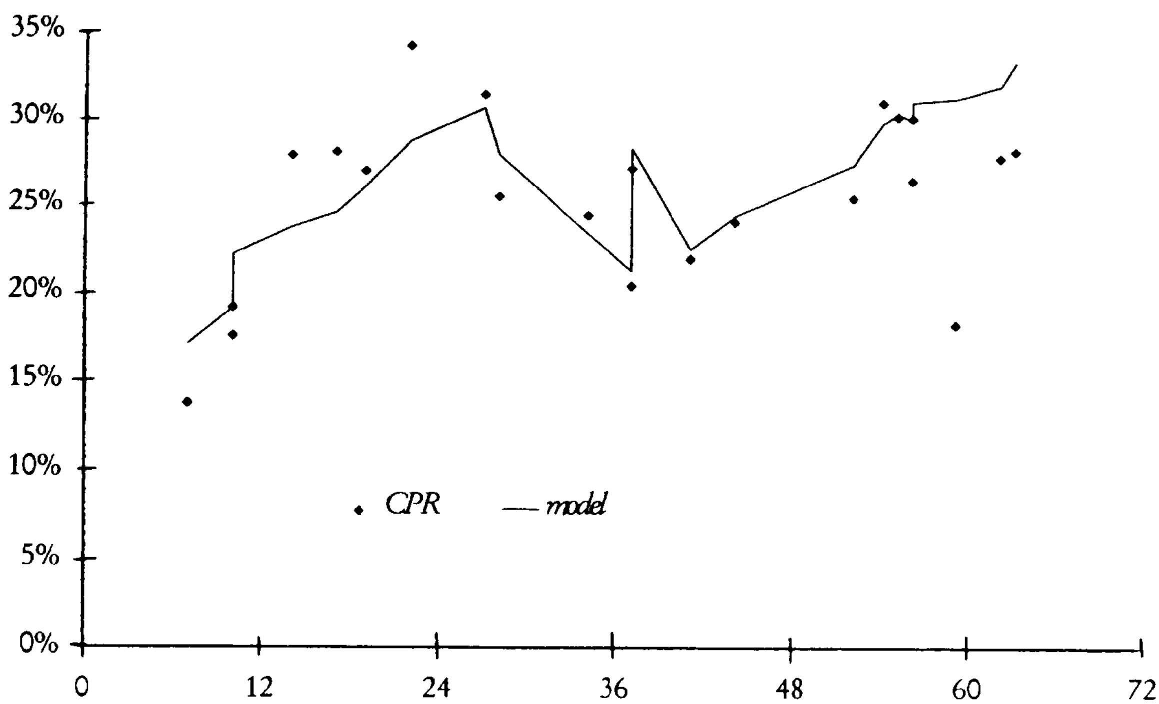

To give you a couple of examples of how much difference the optionality makes, I've plotted a couple of graphs. First, for a pool of 15-year loans, home equity loans aimed at B and C borrowers, you would not expect these to be too interest-rate-sensitive, but you expect then to be somewhat sensitive. When you vary the interest rate scenario, the average life on that pool varies in the bank of about one-and-a-half years (Chart 7). This is more or less what you would expect. The optionality is not too drastic, but it's certainly not something that you'd want to ignore completely. That's borne out by looking at our price yield curve for that security. The dotted line in Chart 8 shows you the price yield curve that you get if you ignore the optionality on the security. The solid line shows you the price yield curve you get from using the full prepayment model. For high levels of interest rates, there's really not much difference, but for lower interest rate scenarios and, in particular, for the kind of scenario that we have now, that is a fairly appreciable difference between using a model and not using a model.

FROM THE FLOOR: In what way has your model differed from actual prepayments recently?

DR. PHOA: Prepayments have been somewhat quicker than our model would have predicted, and the reasons for that are kind of interesting. What we saw at the end of 1997 and early 1998 was that there was a lot of consolidation in the home equity loan business. Green Tree was purchased, for example. A couple of big lenders consolidated, and the actual market has been shrinking. The market could have gone in two different directions. On the one hand, you could see refinancings go down because there was less competition in the market. On the other hand, you could see refinancings accelerate because if you have bigger players in the market, they might be more aggressive in their pricing because of economies of scale. What we've seen is the second example in which case prepayments have really been quicker than anyone would have predicted at the beginning of the year, because consolidation has probably driven down margins rather than shut off competition. But that would have been a difficult call to make 12 months ago.

CHART 7
Pool Average Life Sensitivity
New Pool of 15-Year Loans, 11% WAC, HEL_SP2 Parameter Set

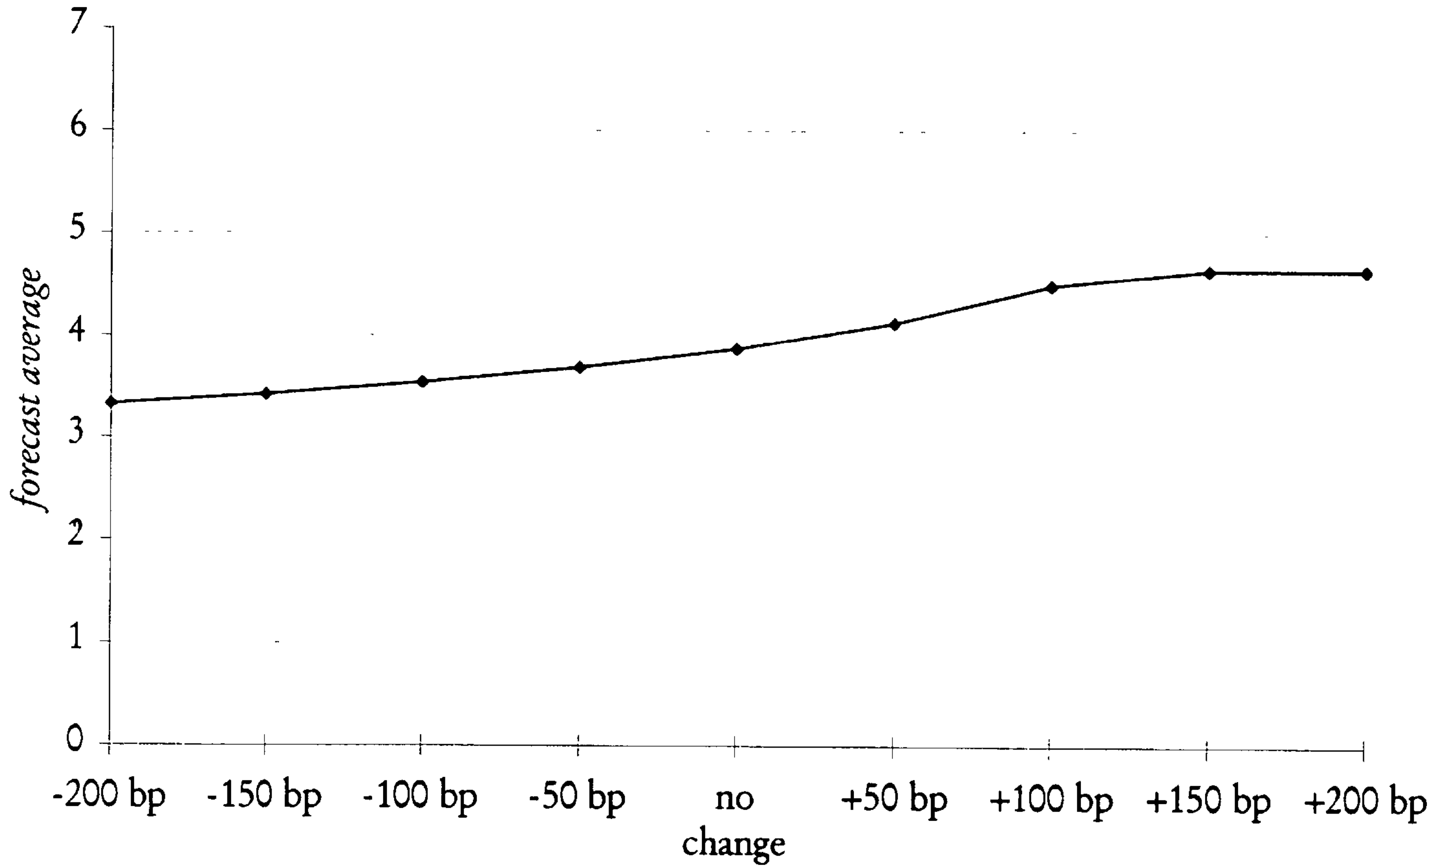

CHART 8
Pool Price/Yield Curve
New Pool of 15-Year Loans, 11% WAC, HEL_SP2 Parameter Set

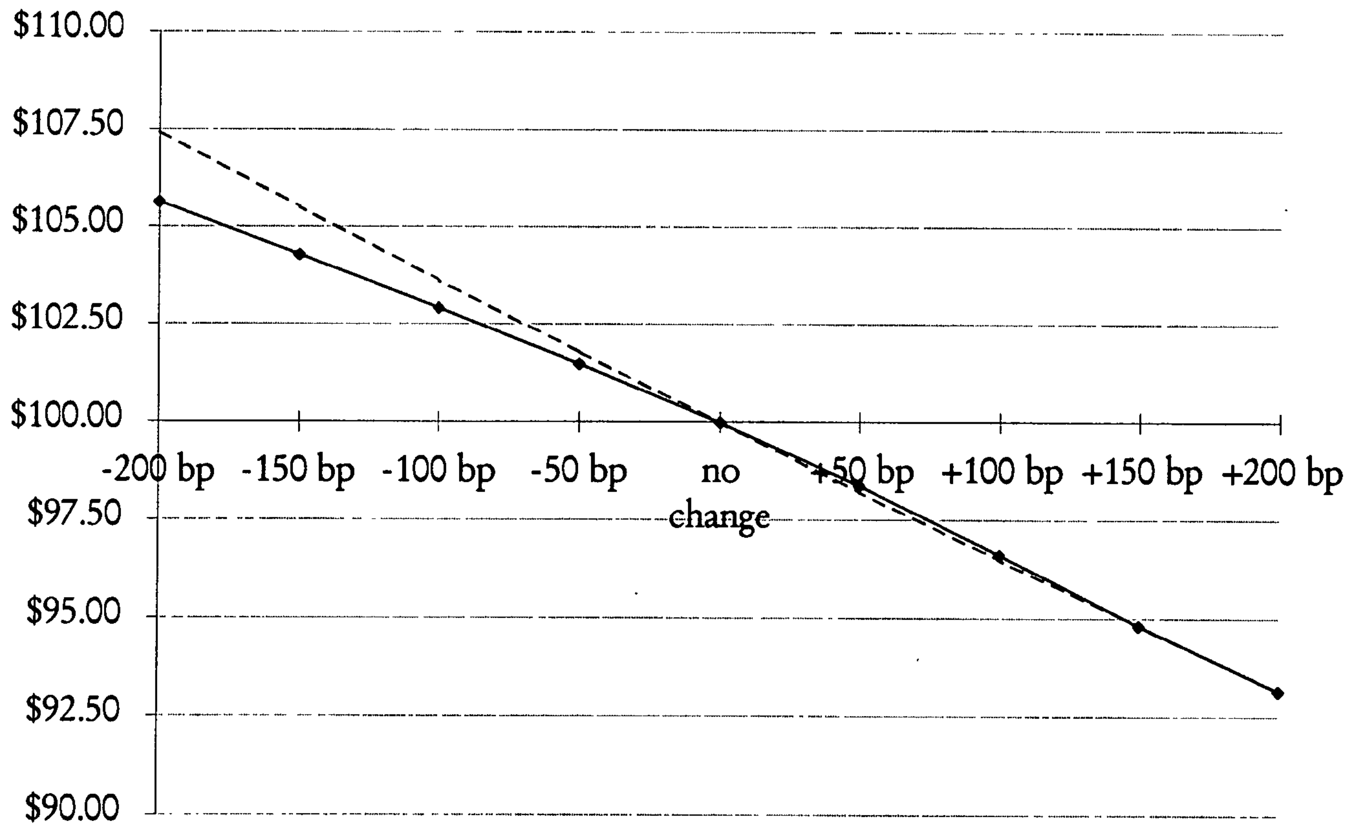

I've just brought up some screens from the BondEdge product just as an example of the kinds of analytical tools you'd use to analyze securities like this. They're essentially the same tools that you'd use to analyze any mortgage-backed security or collateralized mortgage obligation (CMO) First is information from the security description screen. Table 1 is a summary of the deal. You can run off cash-flows under different scenarios. Here's a static cash-flow profile. You can see that this security has a jagged-looking cash-flow profile, and that's because you tend to find in pools of home equity loans and in manufactured housing pools that the collateral itself tends to be more heterogeneous than in pools of conventional mortgages. So, modeling an asset-backed security introduces quite different kinds of problems from modeling a CMO. When modeling a CMO, you need to get the payment rules correct, and the payment rules are very complicated. The behavior of the security is very sensitive to the payment rules. In a typical asset-backed security deal, the payment rules are very simple, but the collateral is very complicated, and you need to focus on modeling the collateral correctly. Modeling asset-backed securities (ABSs) is very similar, in principle, to modeling CMOs; on a practical level the focus of attention and where you devote your resources are quite different.

MR. JACKSON: Is that bond that you showed in the security calculator, modeled by CMS, and is that part of the local database that individuals would use or does CMS model that particular security in its database?

DR. PHOA: CMS models this security. At the moment we try to provide complete coverage of agency and whole loan CMOs, home equity loans, and manufactured housing. Up until early this year we did have a gap. We didn't have home equity lines of credit on the database, and we've recently added those deals as well. We're in the process of adding commercial mortgage-backed securities. That bond is modeled in the CMS database which is available to all of our users. We are committed to keeping up with issuance of these kinds of securities.

TABLE 1
Analysis of a HEL Security

Security Calculator - Asset-Backed Pricing Date: 08/27/98	
CUSIP	21075WFJ
Issuer	CONTIMTG_HEL_1997-03-A9
Class #	9
Moody's Quality	AAA
Sector	ABS
Coupon Rate	7.120
Maturity Date	08/15/28
First Principal Payment	07/15/97
Payment Frequency	M
Issue Date	06/15/97
Structure	PASS
Cleanup Call(%)	0

Collateral		
Orig Size ($MLN)	945.000	
Type	MIXED	View Call
Wtd. Avg. APR(%)	MIXED	
Wtd. Avg. Maturity (Mos)	MIXED	
Prepayment Assumption(%)	129	PPC
Current Price		104.736

ABS Deal Summary

Deal	Original Collateral
CONTIMTG_HEL_1997-03-A9 Original Size ($MM): 945.000 First Pymt (MM/YY): 07/97 Number of Classes: 13	Type: MIXED

Tr #	CUSIP	ID	Type	Coupon	Maturity	Initial Balance	Remaining Balance
1	21075WFA	A1	SEQ	6.420	04/15/07	130.000	0.000
2	21075WFB	A2	SEQ	6.510	05/15/12	135.000	19.089
3	21075WFC	A3	SEQ	6.680	05/15/12	215.000	201.602
4	21075WFD	A4	SEQ	6.820	05/15/12	73.000	73.000
5	21075WFE	A5	SEQ	7.010	08/15/13	71.000	71.000
6	21075WFF	A6	SEQ	7.130	03/15/15	36.000	36.000
7	21075WFG	A7	SEQ	7.280	05/15/24	65.000	65.000
8	21075WFH	A8	SEQ	7.580	08/15/28	38.600	38.600
9	21075WFJ	A9	SEQ	7.120	08/15/28	68.000	68.000
10	21075WFS	A111	IO	8.500	12/15/99	0.000	0.000
11	21075WFK	M1F	SEQ	7.310	08/15/28	54.337	54.337

Of course, that's not possible for private placements. As I was explaining, the important part of modeling an asset-backed security is modeling the collateral. If we have a security where we can't get any information at all on the collateral, then it's clearly impossible for us to put on the database In that case, we don't try because, in our view, if we can't model a bond accurately, then it's better for us to work with each individual client to try to model it with the information they have than to try to put something inaccurate on the database. That can be quite dangerous.

The other screens give some examples of simulated returns (Chart 9 and Table 2). An important point to make here is that you're not just getting a different return under different scenarios. You can see that your duration is drifting around as well. If you're trying to immunize liabilities or track some kind of benchmark, then you have to take that duration drift into account, and that arises mainly because of the optionality. Then you can run various sorts of cash-flow testing simulations. I've just shown our cash-flow testing screen. That's forecasting not just cash-flows on the security but also book value that you can't see. Book yield and other statistics like that are off on the right-hand side of the screen. In the second half of this presentation, I'm going to say a little bit more about these kinds of tools anyway.

There are four main points I wanted to make. First, prepayment modeling is harder than you think. At one stage it was very tempting to believe that if only you had enough data, and if only you had enough computing resources, and you're a good enough statistician or you're a good enough econometrician, you could get a prepayment model that really pinned down what was going to happen. That was perfectly reliable. It has become pretty obvious to everybody—I guess it was always obvious—that that's really impossible. There aren't any perfect prepayment models There's absolutely no way to get rid of model risk. In another sense, prepayment modeling is easier than we thought it was because, although the econometrics is important, although all that statistical estimation does play an important role in estimating models, it's not quite as central as we thought. In fact, the experience that we all have, and the common sense, intuitions that we all have about what makes people prepay their mortgages are really just as important, particularly when you're looking at new asset classes for which there isn't that much historical data. In that sense anybody who's

CHART 9
Static Cash-Flow Profile
21075WFJ—Cash-Flow
Calendar (08/27/98–08/15/28)

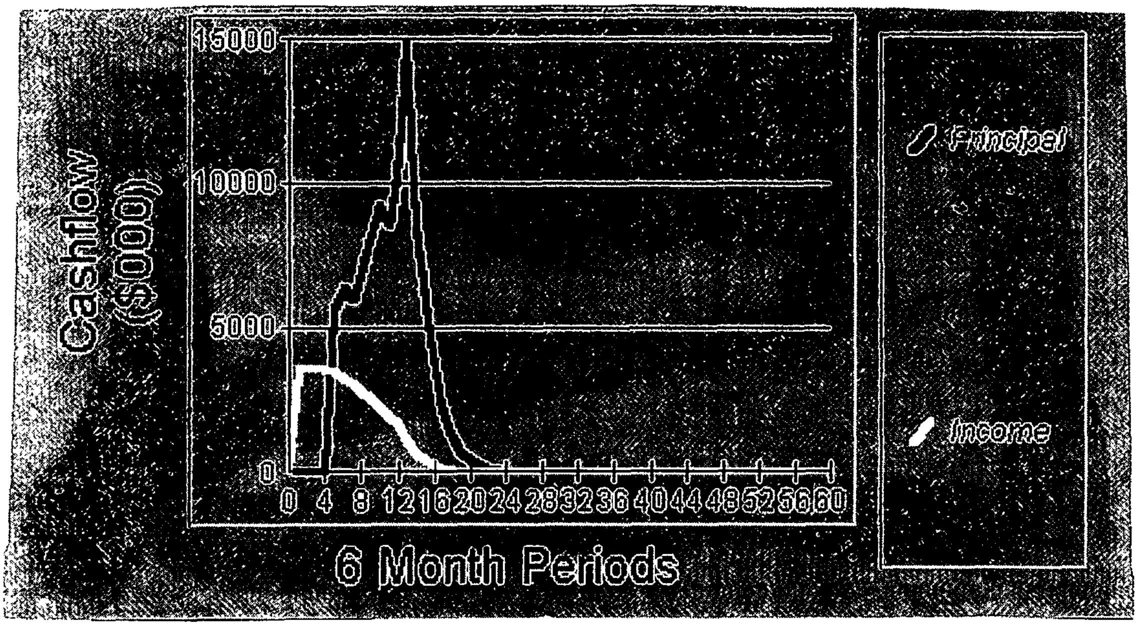

prepared to devote a lot of time to analyzing the market and looking at market conditions can have some useful insights into prepayment analysis. The key here is to use all the data that are available and to use them efficiently and in an intelligent manner, which means you don't just feed it into some statistical machine. You try to interpret it and see what it means in terms of your understanding of the mortgage market.

The fourth point is really the most important one: once you have built a model you still need to focus very carefully on model risk. You need to think about the different assumptions behind your model, how they could go wrong, and how your model could drift away from reality. That's the most important message of prepayment analysis.

MR. JACKSON: I think we'll just start right in on the next part of the session. The only other comment I'd make was that CMS is not alone in seeing the assumptions drift away. At my firm we've looked at Salomon Brothers yield book assumptions for quite a while. Salomon Brothers

TABLE 2
Horizon Return Simulation

Single Bond–Total Return   Pricing Date: 07/31/98

CMS

Horizon Months	3	**Issuer:** CONTIMTG_HEL_1997-03-A9	
Reinv. Rate	5.110	**CUSIP:** 21075WFJ	**Mod Dur:** 4.31
Increments (Bps)	50	**Coupon:** 7.120	**Eff Dur:** 4.20
		Maturity: 08/15/28	**Conv:** 0.02
		Yield: 6.497	

				--- CPR (%) ---		
Yld Chg	**Return (%)**	**Price ($)**	**Eff Dur**	**Lifetime**	**12 Month**	**PPC**
-300	14.25	115.801	3.84	32.91	32.91	164.53
-250	12.09	113.572	3.88	31.86	31.86	159.31
-200	9.95	111.367	3.90	30.78	30.78	153.91
-150	7.86	109.209	3.92	29.64	29.64	148.18
-100	5.80	107.089	3.93	28.42	28.42	142.11
-50	3.77	104.993	3.96	27.14	27.14	135.72
0	1.72	102.880	4.01	25.80	25.80	129.00
50	-0.27	100.830	4.08	24.39	24.39	121.93
100	-2.30	98.740	4.17	22.89	22.89	114.43
150	-4.35	96.631	4.27	21.29	21.29	106.47
200	-6.41	94.514	4.37	19.65	19.65	98.25
250	-8.46	92.409	4.46	17.99	17.99	89.97
300	-10.50	90.314	4.55	16.34	16.34	81.68

periodically changes their assumptions as well. Nobody has the right prepayment model. CMS is not alone in this drift-away concept.

DR. PHOA: Now I'm going to talk about managing option risk in a portfolio. Let me just start off by saying that this talk was originally written for a slightly different audience. Please bear that in mind. I think that I look at these problems from a slightly different viewpoint than most of you because my own background is on a trading desk for Deutsche Bank. It's a slightly different perspective than an actuary's. Please make some allowances for that.

I'm going to run through a couple of definitions of option risk and look at the kinds of securities that are affected. I'll say something about the impact of option risk and give a few examples, and then spend some time talking about different models like term structure models and prepayment models. I'll try to make some kind of comparison between different approaches towards modeling option risk. I've left out most of the mathematics. You an get hold of that from me if you're really interested, but I think it's more important in a talk like this to focus on the important intuitions.

I'll begin by giving a couple of different definitions of option risk, all of which are slightly different, though more or less equivalent. The first, which really is the commonest, is just that securities have option risk if the size or the timing of the cash-flows on those securities depends on the level of interest rates. That's the simplest definition. In some way, all the others follow from that.

The second definition is that the return on those securities is a nonlinear function of interest rates. You can see that—except for the fact that even an ordinary bond has some convexity—it kind of follows from the first one. The third definition captures just a part of option risk but an interesting part which is that the value of those securities has some exposure to implied volatility in the over-the-counter option markets. That's a function of the fact that volatility expectations affect the weight that people attach to different interest rate scenarios in the future. When the cash-flows are sensitive to the level of interest rates, that affects the way securities are marked-to-market.

A very important point to focus on is that option risk isn't a function of accounting. Even if you have assets that are held at their book value, they still have option risk. It's just realized in a different way. That's an obvious point, but it's one that's worth repeating.

Where does option risk come from? It can come from different places. It can come from an investment portfolio. You might have callable bonds, mortgage-backed securities, or even home equity loans, as I was saying before. It can come on the liability side. For example, if you're a corporate issuer, and you issue bonds with put options, then you have option risk on the liability side. There are many more complicated ways in which insurance-related liabilities can have option risk. There are various business sources of option risk. For example, if you're a corporation that relies on short-term funding, then your access to shorter-term facilities might be affected by certain covenants if interest rates rise sharply. There's also the whole area of real options, which I won't go into.

Let me give just a few quantitative examples of the impact of option risk in practice. The first example is a callable agency bond (Illustration 1). It's a 10-year bond. It's callable in two years' time and thereafter at par. The optionality is quite apparent. The modified duration of the security is quite long, i.e., that of a 10-year bond, but it has a much shorter effective duration. It has a negative convexity, and the actual value of the option in this case is quite large, or nearly five points. From a mark-to-market point of view, I've prepared a total return simulation. The option risk on the security is very apparent. In a high interest rate scenario, it's really behaving like a 10-year bond, and in a low-interest-rate scenario, for obvious reasons, it's really behaving like a two-year bond.

For more complex securities like CMOs, the place to start in analyzing any kind of securities option risk is to look at graphs of return simulations, and to look at graphs of cash-flow profiles under different scenarios because the first step is always to get an intuition for how the security behaves and then to start doing the quantitative work.

Let me illustrate that with a Fannie Mae mortgage. At the time that I ran this analysis, it was a pretty cuspy one. Now it's at quite a premium. It's just a pool of 7.5% 30-year mortgages, issued in 1996.

Chart 10 shows a pattern of cash-flows under three different interest rate scenarios. The solid line shows the cash-flows under a scenario with no-changes. They're just declining over time because the actual pool or number of mortgages is declining over time. The dotted line shows the cash-flows under an up-100-basis-point scenario. There are fewer cash-flows earlier on, but, of course, later on the cash-flows are higher because the pool balance is higher. The opposite is the case for the thin solid line which is a down-100-basis-point scenario. Initially you're getting higher cash-flows. In 10 years' time, you're getting lower cash-flows because there are hardly any mortgages left in the pool.

In what different ways does option risk affect you then? It changes your cash-flow profile. Your cash-flows extend or contract depending on what happens to interest rates. It changes to a horizon return. It doesn't just affect cash-flows. It also affects mark-to-market. What's equally important is it also affects your risk profile. For example, the duration of your portfolio can vary quite drastically if interest rates fluctuate. As well as getting an impact on your mark-to-market, you also get a tracking error versus whatever it is you're trying to match off. As a rule-of-thumb, your duration always moves the wrong way. When the market rallies, you always get shorter. When the market sells off, you get longer. That's what negative convexity means.

Table 3 shows some more horizon returns on that mortgage pool. You can see the reason for the optionality: the prepayment speeds under different scenarios, vary drastically as interest rates move up and down and that's having an effect on your returns and on your duration. In this case, the duration of that mortgage under different scenarios varies between 1.5 and 5.5, and the actual relationship between duration and the level of interest rates can be very complex. What that means in practice, of course, is that if you're trying to track a benchmark or if you're trying to match off liabilities, in order to stay matched, you have to continually be rebalancing your portfolio as its duration drifts around. From a trader's point of view, that's called gamma hedging. You have a negative gamma position, and that always costs money.

ILLUSTRATION 1
Callable Agency Example

Security Calculator—Bullet Bond Pricing Date 08/31/98

LOCAL

CUSIP	10NC2
Issuer	TMA Test
Moody's Quality	AGY ▼
Sector	AGY ▼
Coupon Rate	6.000
Next Coupon Date	2/15
Maturity Date	08/15/08
Coupon Frequency	SEMI-ANNUAL ▼
Issue Date	00/00/00

Pricing Method

- ● CMS
- ○ Spread to Treasury
- ○ Manual Price
- ○ Manual Yield
- ○ IDC AutoPrice

Options

Call Price	100.00	Sinking Fund:	
Call Date	08/15/00	% Ret/Yr.	0.00
Put Price	0.000	Start Date	00/00/00
Put Date	00/00/00	End Date	00/00/00

Current Price	100.168
Yield to Maturity	5.977
Yield to Call	5.906
Yield to Put	0.000
Current Yield	5.990
Modified Duration	7.403
Effective Duration	3.718
Convexity	-0.747
Option Value	4.995
Accrued Interest	0.267

CHART 10
Multiple Scenario Cash-Flow Analysis

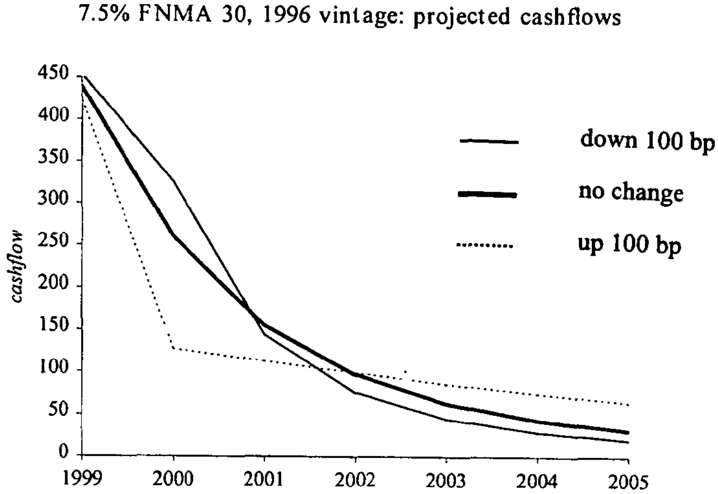

Let me spend a little time talking about some models, and then I'll try and wrap up the presentation. If you're analyzing securities of option risk, there are basically two classes of models that you need to build. First, you need to build term structure models. They're the models that you use to analyze both callable bonds and mortgage-backed securities. Second, you need to build prepayment models, and they're the models that you use to analyze mortgage-backed securities specifically.

I've already spoken about prepayment models, so let me focus more on the term structure side. What is a term structure model? It's just something that describes the random changes in interest rates. In other words, given today's term structure, it gives you a description of the probability distribution of different interest rate scenarios at various times in the future. Although there has been a tremendous amount of work on different term structure models, you would not see many different term structure models commonly used. There's maybe half a dozen variations. Not always, but in most cases, they give broadly similar answers for different securities, so there's a reasonable level of agreement. That's completely different from the situation with prepayment models where many people have many quite different prepayment models. Sometimes you can see them giving wildly inconsistent answers for different securities. All you have to do is take a single CMO and compare the option-adjusted spreads that different brokers compute for that CMO to see that.

What are the issues that crop up when you're trying to choose a term structure model? I'm assuming that most of you will not be actually getting down and building and implementing your own term structure model. What you're really faced with is making a choice, choosing one or two term structure models that are being offered by different vendors and that might mean vendors who are offering a complete package like we do or that might mean vendors that are offering components. In other words, what you're trying to do is make a decision about a number of different issues. I think there are three important issues in assessing a term structure model. First, what kind of interest rate process does it use? In other words, what's the shape of the probability distribution of outcomes? Second, does it have mean reversion built into it? Does it assume interest rates follow a random walk or does it assume that they do tend to get attracted back to some long-run average? Finally, it is a single-factor model or a multi-factor model? Those aren't the only issues, but I think they're three of the most interesting ones.

To clarify, what does it mean to choose between different interest rate processes? That's a choice between different probability distributions. The distribution of interest rate shifts might be normal, or lognormal, or they might be somewhere in between. An example of what that means would be if you're assuming that right now there's a certain basis point volatility in yield, does your model assume that if yields move up or down that your basis point volatility is a constant? Then it's a Gaussian model. Or does it say that basis point volatility moves up and down proportionally to the level of yields, in which case it's a lognormal model?

Table 4 shows four different models. Under a no-change scenario that might assume a current 60-basis-point annual volatility. The Gaussian model assumes that in an up-or-down shift scenario you also have a volatility of 60 basis point. However, the other models assume that in up-shift scenarios, you have higher basis point volatility. In down-shift scenarios, you have a lower basis point volatility. That's obviously going to affect the results you get to some extent.

Chart 11 compares the results that you get from assuming different kinds of processes when you're evaluating a single callable bond. The basic message here is that using a Gaussian model or using a normal model makes callable bonds look more risky than using a lognormal model. The reason

TABLE 3
Horizon Returns

Single Bond–Total Return Pricing Date: 08/31/98

CMS

Horizon Months	3	Issuer:	FNMA
Reinv. Rate	4.989	CUSIP:	FN075026
Increments (Bps)	50	Coupon:	7.500
		Maturity:	06/01/26
		Yield:	6.350
		Mod Dur:	2.44
		Eff Dur:	1.98
		Conv:	-0.72

Yld Chg	Return (%)	Price ($)	Eff Dur	— PSA speed (%) — Lifetime	12 Month	CPR
-300	5.29	108.618	2.10	644	985	38.64
-250	4.51	107.501	2.05	644	985	38.64
-200	3.73	106.402	1.92	644	984	38.64
-150	3.00	105.354	1.68	641	980	38.43
-100	2.42	104.433	1.53	637	975	38.22
-50	2.04	103.647	1.65	575	847	34.49
0	1.47	102.681	2.12	466	636	27.96
50	0.33	101.221	2.91	327	386	19.61
100	-1.38	99.269	3.78	207	188	12.42
150	-3.55	96.929	4.49	149	111	8.91
200	-5.96	94.372	4.98	130	104	7.80
250	-8.45	91.762	5.27	120	99	7.19
300	-10.95	89.139	5.45	111	88	6.66

TABLE 4
Interest Rate Process

	-200 bp	-100 bp	NC	+100 bp	+200 bp
Gaussian (normal)	60 bp	60 bp	60 bp	60 bp	60 bp
CIR (square root)	49 bp	55 bp	60 bp	65 bp	69 bp
Lognormal	40 bp	50 bp	60 bp	70 bp	80 bp
Constantinides-Ingersoll	25 bp	46 bp	60 bp	76 bp	92 bp

is that a call option becomes more valuable when interest rates fall, and it's the Gaussian model that assumes that volatility stays where it was. It's the lognormal model that assumes that basis point volatility falls. The lognormal model makes a callable bond look less risky than a normal model. Similarly, a lognormal model makes a putable bond look more attractive than a callable bond because, when interest rates go up, it's the normal or the Gaussian model that assumes that basis point volatility stays the same. However, it's the lognormal model that assumes the basis point volatility goes up proportionally. In other words, lognormal models make callable bonds look better, and they also make putable bonds look better. I've always suspected that's why brokers prefer to use lognormal models when they're analyzing bonds.

The second issue here was mean reversion, and this is a question about whether interest rates follow a random walk; in other words, the shifts in interest rates tend to be persistent. If you like to think in terms of time series, you have a unit root or interest rates that tend to revert to some long-run average, which is known as mean reversion. It's hard to decide what sort of process interest rates really follow. The main point here is that mean reversion does have a lot of desirable properties, and it seems like it's an attractive thing to incorporate in a term structure model. Why is that? There are three groups of reasons why it's a good feature to have, although none of them are completely decisive.

CHART 11
What Difference Does It Make?

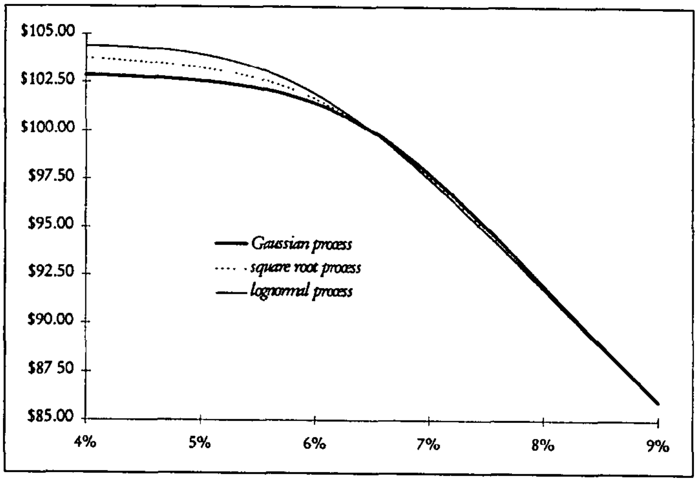

The first is a technical reason: if you don't have mean reversion in a term structure model, then in order for it to be arbitrage-free, the model implies that the yield curve has to steepen drastically. For example, if you specify that yields always move in parallel, and there is no mean reversion, you can prove mathematically that the yield must on average steepen by 400–500 basis points over the next 30 years. That's a strange conclusion that the yield curve has to steepen drastically or systematically over time. It's unpleasant, and it's something that you can avoid by putting mean reversion into your model. That's just a purely technical reason why you want mean reversion. It's to stop your model coming up with unintuitive behavior that forces making it arbitrage-free.

The second reason is empirical. Mean reversion is consistent with what you observe in the markets. It implies that long bond yields are less volatile than short bond yields. Now, that hasn't necessarily been true in the past 12 months, but it's true generally. Mean reversion is consistent with the empirical data in that sense. From a practical point of view, it seems as if mean reversion is more consistent with the way that long-dated options are actually priced in the market. That's slightly circular reasoning, of course, but when combined with the others, it has some weight.

The third decision you have to make is between using a single-factor model and using a multi-factor model. Put questions of accuracy to one side for the moment. There is a very, very compelling reason to use a single-factor model rather than, say, a two-factor model which is much quicker to run. At CMS we've experimented with implementing various kinds of single-factor models and two-factor models. It is very hard to get a two-factor model that runs less than 15 times as slowly as a single-factor model, and if you're analyzing securities individually, that's not too bad, but if you're analyzing a portfolio of 1,000 bonds, then a factor of 15 in efficiency is pretty painful. The only possible way to get that down seems to be to accept a big compromise in accuracy, which kind of defeats the purpose anyway.

There are certain kinds of securities, such as certain sorts of structured notes, where it does seem important to use multi-factor models. The general conclusion, which is a little bit surprising, is that for most securities, like most callable bonds or most mortgage-backed securities, there isn't that much difference between using a one-factor model (as long as it's a good one) and using a two-factor model. Chart 12 shows how the two models are not giving you identical answers, but they're giving you surprisingly close answers, and the difference between these two lines is much less, for example, than the difference between using two different prepayment models. Except in special cases, you aren't really losing that much from using a single-factor model rather than a two-factor model, and there are theoretical reasons for that.

CHART 12
Single Factor vs. Multi Factor

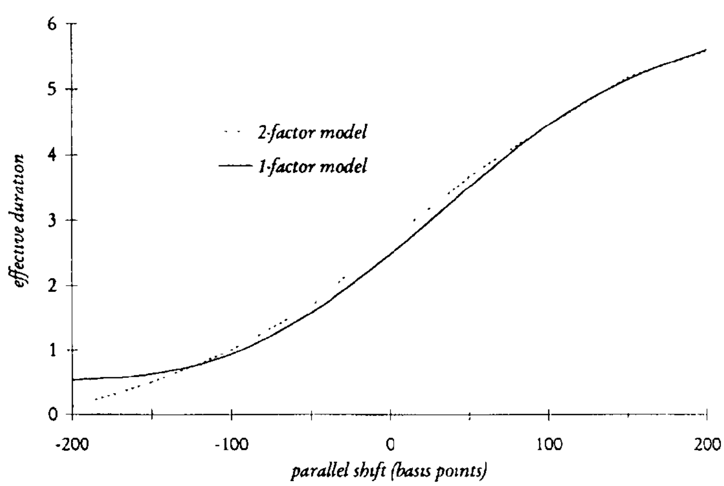

I'd like to make a few remarks on implementing models and on experimenting with different kinds of models. First some remarks on implementation. The main point here is that once you actually have a term structure model, and once you actually have a prepayment model, then you actually have to implement it. You have to put into place some kind of Monte Carlo method for generating random interest rate paths to evaluate a CMO or an asset-backed security. There are many, many different methods of generating random paths. Some are extremely inefficient, and some are extremely efficient. They're all tuned for different situations. If you're comparing two implementations, then just because one uses 200 paths and the other one also uses 200 paths, that tells you nothing about the relative accuracy of those two. If one uses 200, and the other uses 30, that doesn't mean that the one that uses 200 is necessarily more accurate. You have to look at how those paths are generated, and it could well be that the one that's using fewer paths is giving you a more reliable answer because it's using a smarter way of generating paths.

The path generation method that's used in the BondEdge product is tuned to give you a very good, accurate answer for fixed-rate securities, using only a relatively small number of paths. With 16 or 32 paths, you can get an answer that's as accurate as if you generated several hundred or even a thousand Monte Carlo paths. And that's something that you can see for yourself empirically by adjusting the number of paths and looking at convergence. It's based on a series of optimizations that don't seem that impressive but that combine in a very nice way. It's specifically tuned for fixed-rate securities, but it does have limitations. You need to run more paths for adjustable rate mortgages with tight reset caps. It's important to bear in mind that efficiency depends on the implementation, and there really are a lot of different ways to implement models like this once you specify them.

I just wanted to say a little bit about some new lines of research that we've been following at CMS, and this is not so much in the direction of building prepayment models for use in production. It is meant to try to get a handle on the cause of the prepayments in much greater detail and get a handle on model risk and quantify it in much greater detail. Instead of taking an econometric approach to modeling prepayments and trying to estimate different response functions from historical data, we've attempted to go to a new level of detail and explicitly model the decision rules that borrowers use and to explicitly model the composition of pools over time. For example, a borrower might say, "I'll refinance my mortgage if I can save $2,000 in my monthly payments after transaction costs." The advantage of that is you can incorporate explicitly the effect of changing fixed and variable transaction costs, and you can incorporate explicitly the impact of the expected time to relocation that a particular borrower has.

This turns out not to be a really practical tool at a production level because models like this, because they're so detailed, are extremely slow. Models like this are useful for analyzing model risk. For example, given what your prepayment model predicts now, then models like this will tell you what happens if transaction costs fall by $700. These models indicate what happens if closing underwriting standards change and the average points paid for a refinanced mortgage rise from one point to two points? Models like this let you make an explicit and quantitative link between things that happen in the mortgage market and quantitative prepayment estimates.

FROM THE FLOOR: In your comparison between one-factor and two-factor models, which models were you using?

DR. PHOA: Let me explain precisely what these two models are. The one-factor model is a CMS model. That's a normal model with mean reversion. In the two-factor model, the two factors are changes in the level of the slope of the yield. They are changes in the level of the yield curve and changes in the slope of the yield curve. Innovations in both of those series were assumed to be normally distributed. The one-factor model assumes not-quite-parallel shifts in yields. The two-factor model assumes two kinds of yield curve shifts: parallel, and slope shifts. If you look at the range of different yield curve models, it seems as if the most realistic ones are the ones that take as their two factors parallel and slope shifts. If you do a principal components analysis on an historical time series of bond yields, you see pretty clearly that parallel and slope shifts are the most important factors that do explain most of the variation in bond yields. Those are the assumptions that are made in that model.

There are other choices you could make for the second factor. You could look at changes in the volatility of the short rate. That's the Longstaff and Schwartz model. There have been various criticisms of different two-factor models that I won't go into. Broadly speaking, in other comparisons of one-factor versus two-factor models that I've seen you get similar results. Except in special cases, a two-factor model does not necessarily add that much reliability or accuracy. And in the case of mortgages, your time is better spent improving the prepayment model, as that's likely to be a larger source of error.

Similarly, at some point, there's no point in refining your implementation of the Monte Carlo simulation because once you run up against the inherent inaccuracy of any prepayment model, and once you run up against the inherent limitations in model risk in your prepayment model, then any further accuracy you try to squeeze out of better implementations or out of more accurate term structure models is going to be spurious. It's more important to accept that model risk exists in your prepayment model and to try to gain some control over that. It's true that there are specific examples of securities (adjustable rate mortgages would be an example) in which there's a much bigger difference between the answers that a one-factor and a two-factor model give.

FROM THE FLOOR: Wouldn't a two-factor model be necessary for modeling certain kinds of liability streams, such as insurance policy redemptions?

DR. PHOA: I'm not an expert on insurance liability. I would say it's definitely something with which you could experiment. You just hit the basic trade-off between computation time and richness of analysis.

FROM THE FLOOR: Has CMS considered using low discrepancy sequences to generate interest rate paths?

DR. PHOA: The method that CMS uses is proprietary, but what it actually corresponds to is a version of the method of good lattice points, which has been pretty well studied in the quasi-Monte Carlo literature. Maybe I can give a bit of background there.

The Monte Carlo method, which generates paths purely at random, is very inefficient. Accuracy only goes down as the square root of a number of paths. So, various methods have been devised to generate paths to give more uniform coverage and have better convergence. The two main approaches to improving our Monte Carlo paths have been, first, low discrepancy sequences such as Sobol sequences, Faure sequences, and so on. Second are methods of good lattice points. The main distinction between those is that if you know in advance how many paths you want to use, you use, you use a method of good lattice points, and if you want to keep generating paths until some stopping criterion is satisfied, then you use a low discrepancy sequence. In each of those cases there are many, many different variations. Unfortunately, in practice, the actual efficiency of either of those kinds of methods depends massively on the particular way it's implemented and the particular optimizations that you make. As a general rule-of-thumb, if you want high accuracy, you're probably better off using a low discrepancy sequence because you get more control over when you stop because you can apply a stopping criterion. If you're only going to run a low number of paths, maybe less than a hundred, and you want to specify in advance how many paths you want to run, but you still want good accuracy, then you're probably better off using a method of good lattice points. But it's still a very, very active field of research. We've made a particular decision at CMS. We've experimented, in fact, with using low discrepancy sequences as well. We couldn't quite get good enough results for low numbers of paths, but we haven't given up. It's not in production yet and probably won't be for some time.

FROM THE FLOOR: Which path generation methods work when you have a large number of time points, so that the problem has high dimensionality?

DR. PHOA: It's well-known that in high dimensions low discrepancy sequences can often not work well because the points that you select tend to cluster into a small number of hyperplanes rather than covering the whole space of paths more uniformly. There has actually been a lot of work put into trying to get around that problem of dimensionality, and there are two ways to approach that problem. One is to try to come up with sequences that give you good coverage instead of clustering into hyperplanes. That's very, very hard. Both low discrepancy sequences and methods of good lattice points have this problem of clustering in hyperplanes. The other approach, which is a whole lot easier, is to try to come up with simulation schemes in which you get accurate answers despite the fact that you cluster in hyperplanes. Schemes like that can depend on something other than doing an iterative simulation. You could instead step forward one step at a time to do a recrusive simulation, then jump forward 10 years, and then pick a midpoint. Then, for each pair of points, pick a further midpoint, and so on until you reach a desired level of refinement in time. The order in which you simulate different time points actually makes a big difference to the speed of convergence. That's kind of a subtle question because it depends on the characteristics of the security itself. I'm not going to go into that in too much detail, but you are increasingly seeing research shifting into that area as well as into the area of devising better low-discrepancy sequences.

MR. JACKSON: Wesley, I have one quick question. You said the brokers favor a lognormal distribution process.

DR. PHOA: CMS uses a normal interest rate process. We assume that innovations in interest rates are normally distributed. In other words, our model is giving you more conservative assessments about the risk of a callable bond or of a mortgage-backed security than most broker models. We haven't done that just to be more conservative. That's a choice that's also consistent with what you see in low-interest-rate environments. If you look at the Swiss market or the Japanese market, then there's quite a bit of evidence that, in fact, at lower levels of interest rates basis point volatility gets maintained. It doesn't get driven down. There is some empirical support, although not decisive, for the particular choice that we've made in BondEdge. Of course, I was being a little bit unfair to the brokers because there are practical reasons why we choose a normal model and brokers choose a lognormal model. The other reason that we choose a normal model is that it's much more efficient to implement. That's important if you're trying to run a model on 1,000 different callable bonds. The reason that brokers choose a lognormal model is because it's much easier to calibrate to observe cap and floor prices for technical reasons, and that's important if you have to constantly be recalibrating during the course of a trading day. There are those practical reasons as well as the empirical and theoretical ones.