1999 Valuation Actuary Symposium
September 23–24, 1999
Los Angeles, California

Session 3
Asset Modeling Concepts

Moderator: Frederick W. Jackson
Panelists: Teri Geske†
Charles K. Cackowski‡

Panelists address specific asset modeling issues of current interests. Topics include:

- *Purpose of model considerations*
- *Dynamic reinvestment strategies*
- *Less than rigorous modeling of asset optionality*

MR. FREDERICK W. JACKSON: I was fortunate to recruit two good people. Teri Geske is going to go first, and she's senior vice president of product development at Capital Management Sciences in Los Angeles. I think many of you are familiar with the CMS BondEdge product that's used to support part of the asset modeling work that you do for asset adequacy analysis. CMS, as a firm, is one of the investment modeling group that's more friendly towards insurance companies. She'll be discussing asset model risk, based on a paper that was in the April 1999 *North American Actuarial Journal*.

Ted Cackowski will follow Teri. He's with the Stanford Financial Group in Houston, Texas, and prior to that he was at the University of Texas as an assistant professor of finance for ten years focused on quantitative methods. He's going to follow Teri. She's going to be tying into the asset model era issue, and he's going to go over and discuss portfolio optimization. Part of his presentation is going to be focused on your ability with his processes to examine alternative investment strategies in a timely and cost-effective method. At this point, I'm going to turn it over to Teri.

---

† Ms. Geske, not a member of the sponsoring organization, is President of Product Development of Capital Management Sciences in Los Angeles, California.

‡ Mr. Cackowski, not a member of the sponsoring organization, is with Stanford Financial Group in Houston, Texas.

MS. TERI GESKE: I'd like to make this as interactive as possible. I'm here to talk about model risk in your investment portfolio. I understand a lot of you are involved in determining asset adequacy, and this subject really does have a very direct bearing on judging whether or not you have adequate assets in the portfolio to fund your products. There are really two different kinds of risks in evaluating your investment portfolio; first, there are what I would call market and credit risks. Those are things like the risk of a loss in your investment portfolio due to changes in interest rates or changes in commodity prices or changes in foreign exchange rates or changes in the equity markets. There are also credit risks that are more company specific. There is some danger of default or downgrade with a specific company that has something to do with the credit risk of a particular name, and those are risks that most of us are pretty familiar with. If you've ever attended one of the workshops that my company gives, you'd know we talk at length about duration and convexity and how you measure interest rate risk. So we're all pretty familiar with those two types of risks.

I'm going to talk about the notion of model risk, which is really quite different than interest rate or credit risk. Model risk is a risk that your valuations have some error due to assumptions in your models that are either imprecise, outdated, or they don't hold up during some of the extreme market conditions that you're testing in your sensitivity analyses. So this is a risk that's really quite separate from the market and credit risks that I mentioned above. What I want to point out to you is that model risk is unavoidable. It would be fabulous if I could tell you how to get rid of model risk, but that's not possible. It doesn't mean, however, that you ought to take your models and throw them away, and it doesn't mean that you can't use the output from the models. It just means you have to apply some critical thinking to what the model is telling you.

The concept of model risk, or the concept that the answer your model is telling you may be incorrect, applies to asset valuation. When we value embedded call options, or when we make a prepayment forecast on mortgage-backed securities, or when we look at securities that have embedded caps or floors, or the derivatives themselves, there are some valuation errors that can arise there due to model imprecision. I'm much less familiar with the liability side, so I'll let somebody else speak to that. I know that you have predictions with regards to unscheduled withdrawal assumptions, crediting rates, and expenses that are assumptions. They're not cast in stone from some higher source so your liability modeling is imprecise. By a show of hands, is anybody here using value at risk analysis yet? A couple of people raised their hands. You'll probably you'll be hearing more about value-at-risk in the next year or two. Value-at-risk analyses is very popular among banks. It uses assumptions about correlations across different asset classes and potentially across different types of liabilities to determine your expected loss with a certain probability. It also uses some assumptions about correlations that may be outdated. It uses some historical period of time over which to measure certain risk factors that may or may not be applicable. So if you're using VAR, know that there's model risk in those analyses as well.

Just a fun little example of model risk that I put in here has to do with a kind of investment that was popular in the early 1990s called the yield curve note. In fact, I actually saw one yesterday in a portfolio of a bank that we're working with. I hadn't seen one in a while. The yield curve note had a coupon rate that was based on the spread between two Treasury rates. It would be like the ten-year versus the two-year Treasury or a ten-year Treasury versus a six-month T-bill. The coupon rate would adjust every so often based on the difference between these two rates, and then they would add some spreads. So, in this example, it is the ten-year Treasury, minus the two-year Treasury, plus 450 basis points. I mentioned these were very popular back in the early 1990s. We had a very steep yield curve; about a 300-basis-point difference between the third year Treasury and the one-year Treasury. So the coupon rate based on that very steep yield curve would have been quite high and quite attractive. At the end of 1993, this particular formula would have given me a coupon rate of just over 6%, which would have been about 100 basis points over the ten-year Treasury rate. This is a AAA-rated security, and it didn't really have any credit risk. It was issued by a government agency, so it looked pretty attractive as an investment. What you'd want to ask yourself in valuing this and deciding how much you ought to pay for this security, is what modeling elements that are embedded in my valuation would most affect the timing of my cash flows and/or the amount of my cash flows? Therefore, how would it affect the price of the security or its duration and convexity if you're focusing on interest rate sensitivity. It's all tied together.

For a yield curve note, you really have to see that the amount of your future coupon payments, as this thing adjusts, will depend upon the slope of the yield curve. As the yield curve flattens, that coupon rate is going to get smaller, and you have to realize that the present value of those future coupon payments is affected by the level of interest rates. So you have both the slope of the yield curve and the overall level of the yield curve as key modeling assumptions when you're doing a valuation on a security like this. What's the key model risk for these yield curve notes? Your critical model risk would be a failure to sample a rich enough set of potential yield curve slopes and yield curve levels to adequately capture what your possible future expected cash flows would be from this type of security.

As I mentioned, at the end of 1993, the yield curve was very, very steep. It was 154 basis points between the ten-year and the two-year. At the end of 1994, the slope of the short-term rates was even higher, but over the decade of the 1980s, we would note that the yield curve actually inverted over five of the ten years during the 1980s. So even though you may be living in a particular environment that's attractive with respect to a type of security, history teaches us that environments can change, and you really want to make sure that you sample all these possibilities. Of course, for a yield curve note like this, if the yield curve is inverted, your coupon would drop, in my example, to as low as 4.5% and if interest rates were high (at about 6%), and you have a coupon rate of 4.5%, you're going to be pretty unhappy. The value of that security is going to drop quite a bit. So this is just an example of a particular type of security, and the type of critical thinking you would want to apply when evaluating your model risk.

Now I'm going to focus on the asset side of your asset/liability problem that you all face because that's really the only thing I'm qualified to speak on, but I'm sure that all of you can think of many instances on the liability side of the equation where model risk plays an important role. On the asset side, your model risk affects your valuations and how much should you pay for something or how much you are reporting that it's worth in a marked-to-market sense? How much would you be willing to sell it for? So model risk affects any type of valuation analysis; it affects your risk analysis, such as when you ask yourself, what's the duration of my portfolio, or what's the convexity of my portfolio? Model risk affects those values as well; in fact, some of us are starting to talk about a range of durations instead of a specific duration because sometimes we really don't know what the duration is. And although it would be wonderful to have accuracy out to the fourth decimal point, that really isn't possible. Model risk also affects your forecasts, your cash-flow forecasting, particularly with respect to mortgage-backed securities, which I'm going to talk about in a bit. It affects total return simulations, horizon analyses, and your earnings forecasts. Those are all the things that model risk can affect. If any of you do performance attribution, (for example, evaluating your outside investment managers or in-house investment managers), and you attribute the return on your investment portfolio to certain factors in the market, there's a good deal of model risk embedded in that attribution analysis, and then you have what I call a "mark-to-model." When we're doing simulations or cash-flow forecasts, we're making assumptions about what something will be worth in the future based on average market conditions. When we have extreme movements in the market, as we saw last fall in the fixed income market, those assumptions can completely fall apart. In the fall of 1998, credit spreads moved out by a magnitude that was virtually unheard of. It was a five standard deviation move. That happened over the course of about a week, and all of the correlations that we are all so fond of quoting fell apart and all correlations went to 1.0. That tends to happen under extreme market moves; the correlations all just go to 1.0. All of the beautiful modeling assumptions we have about credit spreads just completely fell apart last fall. So you've got a big model risk when you have a big market move like that and your historical averages and historical correlations don't hold.

So what can we do about this? You can hide under the table; that's always one choice. Let's assume that's not what you want to do. What we would suggest is that you can evaluate or quantify the potential error that an invalid model and assumption would contribute to your answer, and you can decide whether or not that risk is significant. If you tweak your model by 10% and realize that the outcome is really only 0.01 different than what your initial assumption said, then you might conclude that that's okay. There's some uncertainty with respect to the model, but ultimately that particular type of model risk just doesn't matter that much. Alternatively, if you see that by tweaking a key model assumption by 10%, you change your answer by 20%, maybe that's a risk you shouldn't be bearing. That just isn't something you ought to be living with. With respect to evaluating model risk on the asset side, we tweak key modeling assumptions and see how much your valuation changes to help you decide whether or not that model risk is something that you want to assume.

The type of security I'm going to focus on for the rest of my part of the presentation are mortgage-backed securities. This is an area where model risk can be quite significant. Most of you have probably used more than one system for evaluating mortgage-backed securities. You might go to Bloomberg and get a particular screen that says that mortgage-backed security has this prepayment speed and that duration. Then you might go to our system or another system or to a broker on Wall Street and they'll show you the very same security with a different prepayment speed and some different characteristics. You might wonder how in the world these two things can be so different. For the most part, with mortgage-backed securities when you see differences in pricing and in duration/convexity calculations, you can trace it back to the prepayment model. The differences are probably attributable to differences in prepayment modeling, therefore, the prepayment model you use has a big affect on the answer that pops out of your model, whether it's your cash-flow testing or your portfolio simulations or what have you.

No two prepayment models are ever going to produce the same forecast, and you shouldn't expect them to. We all use different data sets to construct our models. We might also use different variables or factors in the models, and some of us might weight each year's worth of data equally. Others might weight more recent years more heavily and assign a lighter, lesser weight to data from early years. These are all very valid things to do; they just produce different answers.

There's also a risk that the prepayment models that we have derived, based on historical data, don't reflect current market conditions. For example, let's say we have a prepayment model that is based on historical data gleaned from refinancing activity over the decade of the 1990s. A big chunk of those data do not reflect the new ability to quickly and easily refinance your home mortgage over the Internet. For most of the decade of the 1990s, it just wasn't possible. So those historical data sets, while they were perfectly valid for their time, might not reflect current market conditions. Even though all of us in the fixed-income business are always updating our

prepayment models, current market conditions are always going to be just a bit ahead of us. Again, the differences in prepayment models means that there's uncertainty associated with valuing and analyzing these securities. That's just something that you kind of have to accept I guess.
Just a brief review for those of you who may not be entirely familiar with what goes into prepayment forecasting. A prepayment model simulates how homeowners will prepay their mortgages. It is based on factors like what we call seasoning or the age of the mortgage, and that really gets at the idea that if I just refinanced my home mortgage two months ago, even if interest rates decline, chances are I'm probably not going to go out and refinance it right away. I'm not going to go through the headache and the paperwork. It also gets to the fact that if I just moved to a new house and, therefore, have a new mortgage, chances are I'm not going to be moving again in the next year so I won't have any prepayment due to housing turnover for very new mortgages. So there's this element of seasoning that goes into prepayment forecasting.

A recent change in the market has affected another critical factor in prepayment modeling, namely how much of a trigger or how much of an incentive people require in order to go out and begin their refinancing process. Years ago, it might have been 150 basis points between what you're currently paying on your mortgage, and what you can get on a new mortgage that might have been required to entice you to go out and do the refinancing. Nowadays, that incentive is more in the neighborhood of 75 basis points or less. We can observe that if someone has a 50–75 basis point incentive to refinance, that is typically the trigger that will start a waive of new refinancings.

The time of year also affects prepayment forecasts. It's not as critical as those first two factors, but we see a lot more in the refinancing activity in the summer months when people are moving because people don't like to take their kids out of school if they can avoid it. There are typically many more refinancings and much more moving during the summer months than during the winter months. There's what we call burn-out, which really refers to people who have had opportunities in the past to refinance, but have not done so. Over the decade of the 1990s, there have been many instances in which interest rates have fallen quite a bit, giving people an opportunity to refinance. If you have someone that took out a mortgage in 1990 and the mortgage is still alive in the year 2000, they've had many opportunities to refinance. Chances are a further decline in interest rates is not going to entice them to refinance. Maybe the value of their house is below the value of their mortgage and they can't refinance, but for whatever reason, we would consider that type of homeowner to be "burned out" if you will. That factor should go into a prepayment model forecast.

One thing I want to emphasize here is that no prepayment model is correct. My company's prepayment model is incorrect. Nobody on Wall Street has a model that's correct. There's no such thing as a prepayment model that's 100% accurate. Hopefully, we do a decent job and we can talk at some point about how you evaluate whether a model is good or not, but there's no such thing as a 100% accurate prepayment model. Just to back up that claim, I'll show you some numbers about what is offered as a dealer median prepayment forecast. It's available from a number of sources. A number of Wall Street dealers contribute their prepayment forecasts on different types of collateral. Table 1 shows the numbers for loans with coupon rates of 6.5%, 7%, 7.5% that originated in 1998, 1996, and 1993. People often focus on that first column, which is the dealer median number. The dealer median says 215% PSA is the forecast for this type of collateral, and if your model doesn't say 215%, therefore, you're wrong. If you look behind that median number 215%, you'll find quite a range of estimates. In this case, estimates range from 151% PSA to 326% PSA, which is over a 200% difference between the high and the low estimate, so we can't get attached to these median estimates. We have to remind ourselves that there's quite a variation beneath these averages.

TABLE 1
Prepayment Speed Estimates: 30-Year Fixed Mortgages
(as of August 1999)

Collateral Type		Dealer Forecast (PSA %)			High vs. Low – Diff	
Coupon Rate	**Originated Year**	**Median**	**Low**	**High**	**Absolute**	**Percent**
6.5	1998	226	175	365	190	209%
6.5	1996	234	176	320	144	182%
6.5	1993	215	151	326	175	216%
7.0	1998	314	213	726	513	341%
7.0	1996	334	235	585	350	249%
7.0	1993	300	219	483	264	221%
7.5	1997	472	344	907	563	264%
7.5	1993	391	311	721	410	232%

For those of you who are not familiar with mortgage-backed securities valuation, let's discuss where the model risk comes in when valuing a mortgage-backed security. The way we model mortgage-backed security is we take the present value of path-dependent cash flows. By path-dependent, I mean that mortgage prepayments are a function of not only where interest rates are today, but where interest rates have been, so it's a path-dependent problem. We take these path-dependent cash flows and we take an average of them over some set of probability weighted interest rate paths. The cash flows are predicted by the prepayment model at each months end along each interest rate path. If the prepayment model is wrong or off by some percentage, say 10%, cash flows predicted in the numerator are going to change. That's obviously going to change the whole answer, the whole number and the whole valuation that you come up with for this mortgage-backed security. The discounted rate in the denominator includes a term "S," the spread over Treasuries, or what we refer to as the option-adjusted spread. That causes the present value of these expected future cash flows to equal the market price of the security today. So if you've ever heard the term option-adjusted spread, and you're not quite sure what that means, you can hopefully glean something from that formula. The R there is the risk-free Treasury rate, to which the OAS is added.

So how do we evaluate the model risk that's due to the uncertainty of our prepayment model forecast? One way is to use a measure that we call prepayment uncertainty. This measure looks at how sensitive our valuation is to a misestimate of the cash flows in the numerator of that formula.

The way we derive this uncertainty measure is we take the model and slow all prepayment speeds down by 10%. When I say *prepayment speeds*, I mean monthly mortality rates that we call single monthly mortality (SMM) in the mortgage modeling business. We don't use a single lifetime PSA estimate. That would be okay, but it's a coarser measure, using SMMs is more precise. We take all the monthly prepayments that our model forecasts along each of these interest rate paths. We slow them down by 10%, and then we increase them above the base case by 10%. We come up with two new prices or values for the security based on these lower prepay estimates and faster prepay estimates. We take the average difference in these valuations versus the base case, and we call this a measure of prepayment uncertainty. Remember that we are not changing interest rates. We're simply adjusting the prepayment forecast up and down 10% to see how sensitive our answer is to an error in our prepayment forecast.

I have a couple of observations. First, the prepayment uncertainty is greater than zero for securities where the underlying collateral is priced above par and, vice versa, the prepayment uncertainty is negative for collateral priced below par. The intuition behind this is that if prepayments are slower than expected, a premium coupon mortgage loan is going to be out there longer than you originally thought it would be. If speeds are faster than you expect, discount securities will be prepaid at par sooner than you expected. So these numbers make sense. The magnitude of the prepayment uncertainty number is related to the price of the security, the underlying weighted average coupon on the security, and the remaining weighted average maturity on the security.

This is the pattern of prepayment uncertainty for fully seasoned, moderately seasoned and new collateral based on the underlying WAC. The y-axis would show these prepayment uncertainty numbers are rather small: -0.1, 0.1, and, at most, 0.5, so at most, a 10% error in prepayment modeling will produce a one-half percent error in the valuation of your security. I'll get back to that since those numbers look kind of small. You might say, "Who cares that they're small?"

You might want to break down that prepayment uncertainty into two components. There's the component where you just adjust a homeowner's response to refinancing opportunities and hold everything else in the model constant (or just your assumption about the rate of housing turnover that's independent of interest-rate-based refinancings) so you can get more finely tuned with respect to prepayment uncertainty. Basically, those two numbers are computed in the same way as the overall prepayment uncertainty measure. You just tweak the model's factor up or down 10% and compute the sensitivity of the valuation to a change in that key modeling parameter.

How can you use this payment uncertainty measure in risk measurement? As I pointed out, these values look kind of small, 0.1%, 0.2%, but remember that we're only tweaking the model by 10%. Remember, the dealer median estimates vary by over 200% in some cases, so if a prepayment uncertainty measure is based on a 10% model tweak, you might want to multiply that by two or three or five or ten to see the magnitude of the potential error. Also, if your prepayment uncertainty number is half the size of your duration, which is the measure of interest rate sensitivity, you can say that a 10% error with respect to prepayment modeling is equivalent to a 50-basis-point move in interest rates, and that's not insignificant. So, you can kind of relate prepayment sensitivity to interest rate sensitivity. I'll point out here that collateralized mortgage obligations (CMOs) can have a much greater prepayment uncertainty than mortgage pass-throughs. That graph I showed you is based on mortgage pass-throughs, and you can have certain types of CMO securities with prepayment uncertainties that are quite a bit larger. I also want to point out that prepayment uncertainty is diversifiable. You can own securities backed by premium collateral and securities backed by discount collateral, thereby offsetting some of your prepayment uncertainty there.

I have a couple of conclusions that I hope Ted will echo. First, risk management is an inexact science. Our conclusion about whether a level of risk is acceptable depends upon the accuracy of your modeling assumptions, and if your potential error with respect to model risk is too high, you may want to restrict the exposure that you have to model risk either on the asset side or the liability side of the balance sheet. Also, if your model risk is small after you've done this type of an analysis, I think you can definitely conclude that it's okay to use more than one system to do evaluation on the asset side versus the liability side. I know a lot of people say, "I wish there was one system that did everything on the whole balance because then I would know all my modeling assumptions are consistent." If you realize that your model risk is fairly small, I think it's okay to use what I would call the best of the breed approach. Use the best system that you can for asset modeling; use the best system that you can for liability modeling, and put the two together. I think you can have confidence in the answer that you derive.

MR. CHARLES K. CACKOWSKI: First, I agree with virtually everything Teri just said, but before I get into the substance of my presentation, I'd like to talk a little bit about what Teri has done and what the ultimate solution might be. In the usual cash-flow testing asset adequacy paradigm, you take some sampling of interest rate scenarios. The paths can be selected to span some range of scenarios to be risk managed or that are stochastically generated, which is consistent with history. At each point in that path, all of the assets, including CMOs, are modeled. Modeling includes both projecting cash flows and determining a market value, given the simulated interest rate environment. For the corresponding liabilities you will, again, have both interest rate sensitive and nonsensitive features. The crediting strategy may be conditioned upon the behavior of assets and liabilities. In a fully dynamic model, each point in the path has to devalue the assets and the liabilities. The strategy should be designed to incorporate as much of the environment as possible, without creating an unacceptable level of model error.

Let's get back to the big picture. Each of the assets and each of the liabilities is going to have a model error at some level. The model errors may or may not be correlated. Model errors are not necessarily mutually canceling or self-diversifying. A strategy should be robust with regard to both model error and structure. A strategy can appear to succeed, not because it is robust under a wide range of scenarios, but because it is exploiting some systematic model error or structure. For example, if a set of scenarios based on a uniform lattice were used, it is possible to construct a strategy to exploit the discrete minimum move. The false structure success can be avoided by using stochastic scenarios instead of deterministic strategies, such as the New York 7.

Vendors such as Bloomberg and CMS provide a plethora of analytical and descriptive data for large universes of securities. Now I want to direct your attention to an article that came out. I just checked the SOA Web site in preparation for this presentation, and I found a quote from Tenny & Gerard in *Risk and Rewards.* They talk about dynamic immunization value. The key concept there is that the strategy contingent cost is measured over the set of all possible dynamic investment strategies. A person trained in math, might start to think, how would I identify all possible random strategies? The answer is, I don't think you can do that right now but, as we'll see at the end of the presentation, we do have some tools that developed in recent years, notably genetic and evolutionary algorithms, which might allow us to examine a subset that is sufficient in some sense. I argue that unrealistic strategies can be excluded by making appropriate assumptions concerning transaction costs. If, for example, volatility is going to cause the strategy to trade excessively, transaction costs would preclude profitability.

Regulatory frameworks are converging. Both bank regulators and the NAIC are asking for transparency and simplicity. There is a section in the 1995 and 1997 draft of the Actuarial Opinion and Memorandum Regulation. I want to focus on the language that deals with documentation of assumptions. I took a quick poll of some of my friends who are actuaries. They tell me whether or not the regulation requires it. The better practice is to provide detailed documentation of the assumptions. For a bond, the call function, the mortgage prepayment function and how you determine spreads must be disclosed. Spreads are used to value bonds on investment and disinvestment. What if you were to ask someone who's in the business of prepayment models for mortgages, "What is your prepayment function for mortgages so I can go ahead and do sensitivity analysis?" What do you think the answer is? Usually "it's proprietary" or "it's not a single function," or "it's a very, very complicated process." If you want to understand the base level of detail, you must create a proxy of your own

I'll now describe the methodology that will allow us to deal with some of these issues. The first step is to take all this detail you get out of your asset analytical vendor and synthesize something that you can compute and compute in a hurry. We've got a piece of software that will run 10,000 iterations, 50 scenarios, all the cash flows, all the liabilities and a dynamic strategy for 20 assets and 20 arbitrary liabilities in about a minute and a half. This really screams if you get compact proxies. Now what does a compact proxy look like? We're going to take a look at two of them. One is mortgage prepayments, and the other is going to be a government agency. One of the tasks of the actuary and the investment analyst is to find efficient proxies, ways that are transparent and make it easy to see model error and then minimize it.

I don't want to hit CMOs too extensively. Bloomberg and a lot of other systems have done an awful lot of work for CMO structure. There is a typical listing of CMO tranches. CMOs take a group of mortgage pools and allocate the principal and interest to anywhere between 10 and 60 tranches (sub-bonds). Vendors, such as Bloomberg, go through all the principal distribution rules, all the floating rate formulas for the interest. They run their prepayment model and produce summary statistics that describe the cash flows. For some of these tranches, we have a prepayment speed versus an average life. Depending on the structure of these, you can have a prepayment speed go from 200 to 400, and it might change the average life of the principal.

For some tranches, it's really horrible. There's a movement of 75 PSA that changes the average life from twenty years to three years. There can be significant model error. What happens as a function of your prepayment model can be volatile.

When Bloomberg runs an option-adjusted spread (OAS), it runs a 32-path model. Effectively, they take a sampling of paths. Let's come up with detailed statistics about cash flows, spreads, and things of that sort. The scenarios they run look very much like the New York 7. They give you cash-flow detail, coupon payments, interest payments, and so on.

We have collected all of this prepayment information from the primary dealers, and Teri discussed what the error looks like. They tell you up front, this is wrong but it's yours. It's expensive too. The constant payment ratio, which is another measure of PSA, is a percentage of the mortgages that prepaid that year. This is the summary of all that detailed financial engineering that they gather from all the experts. Who knows what the computational and historical research effort was in getting that? It's huge, but this is, in some sense, the consensus of what the experts think that collateral is going to do.

I suggest this is the beginning point for you. What are your choices? I have a prepayment model. There's a pretty big error rate in these things. Get a clean proxy that you can pump into your asset/liability, asset-testing adequacy model. Get many iterations and take a look at many models.

All these people are really modeling a couple of things. One is the incentive response and most of them modeled that by what we call an S function or a Sigmoid function. They're taking a look at all the history of these prepayments. One of the things they're trying to measure is how much of an interest rate shift up and down people prepay. They may be doing a lot more, but that's certainly a big part of what they're doing. The other thing they do is they try to understand this concept called age and burn out. What they know is that when you first buy it, you're not going to prepay much of it. That may increase over time. You're sensitive to the aging process, and after a while, it's going to burn out and go down. That looks like the sigmoid function that

you’re all familiar with. It’s a simple parameter function. We have many different people that are trying to estimate when you get down to very computation insensitive areas or very research-oriented processes. They’re trying to estimate one to four parameters on the S function that has shifted, changed its slope, stretched out, or moved. One function lands on the other function in what appears to be a nearly perfect fit.
Now suppose we went to Bloomberg and just for a second grab some scenarios from their complicated kind of jagged graph. Let’s suppose we picked out seven points in time. I’ll call them newer New York seven scenarios that Bloomberg models. These reflect the consensus of what all the smart people have done. It took me just a minute and a half to type these into an Excel spreadsheet from a Bloomberg sheet.

So what we need to do in the modeling business to get to robust, dynamic strategies is to simplify the portfolio by first of all taking the major moving pieces and getting very concise, compact, econometrics measures because I can’t run one custom vector on Bloomberg without a lot of difficulty. Once I get it to this point, I can run 10,000 scenarios in a hurry. So I’m not trying to tell you this is all you need; you need to come up with proxies for lots of things, but always be thinking proxy computational efficiency. You have an assistant that sits above something like BondEdge, sits above something like Tass or PTS. You take those numbers and you reduce them by finding things that you’re really interested in.

I’m going to give you another example of this idea of model error and computational efficiency. That’s going to deal with pricing options and embedded options. We’re going to talk about a callable U.S. agency security. There are a couple of ways you can value options. One is European and another is American. You can take a lattice, a set of pads, and much technology on how you generate these things. You go back to Black Derman Toy. There are many ways you can construct these lattices. Do a lot of computations over every pad and come up with a value for an option

You can also find a closed form in the U.S. such as the Black-Scholes method if that’s modified. The other thing you can do is you can ask yourself, what’s really driving this? For callable U.S. agencies, there are a few simple things. There’s the difference between the long end and the short end of the curve, which is a proxy for risk. Obviously, the longer the lock-out, the more valuable the bond is. Then you can add a random term to it. How good is that? This computation walks along. This is what we use when we do the 10,000 computations, and the 10,000 iterations in a minute and a half. This is simple. Now what do we gain or lose by using something that's a linear proxy instead of a more esoterical lattice formula? The base system or a Black-Scholes modified system? We don't gain or lose much.

What we're really concerned about when you come down to it is the difference between linear and nonlinear. There's the typical textbook graph of strike price versus an option value. It has a little bit of curve to it. If you put a straight line through it, you have an error. There's an error of some sort but it's not huge. Does that matter? What really matters to you is whether or not you're here which is primarily a function of volatility. That's the error you're worried about when you're going to do some computations out on the cash-flow testing process.

If I take my simple linear model and it is a five non-call too, and I take a look at the spread over the last two years, that linear model works pretty well. So when I model this, what I really want to do when I'm trying to do the 10,000 iterations of the 50 scenarios is vary this here because this is a pretty good model as we know. This is the unexpected event in the market. This is the Greenspan speech, this is the thing that I want to allow for the variance increase if I'm worried about model there.

Now let's assume that we come along and we've gotten pretty good proxies for the eight blocks of business that we think replicate what our cash-flow testing does on the liability side. Let's suppose that we reduced our portfolio, even though we have 300 key sips, and some investment person likes to buy a bond every day because that's a good thing to do. Let's suppose that we've gotten into the essence of it and we have a compact number of proxies on the asset side and the liability side. We have some scenario that we want to test. You get back to the Gerard Tenny problem from risk and rewards. What do you do?

All these things are moving at the same time. You have a model on the asset side and the liability side, and you want to know whether or not your strategy is robust. Take mixed reviews in the past, curving large matrices and solving linear programming problems. Grading and service techniques clearly fail. This is more than a national problem. What we're doing is modeling a lot of different things. Depending on what the model parameter looks like, we might model something. It might go up and down. We might model something one time. These are the things that model error might create.

The way we solve this problem is by using genetic and evolutionary optimization techniques. When you're looking at 10,000 iterations, you have to take this on a little bit of faith. That's a whole course in and of itself. What a genetic optimization or an evolutionary strategy will do is it will basically find this line and that line and tell you that's where you ought to be.

I included in the handouts the most laudatory description of evolutionary genetic algorithms I could find with a Web site that will link you to some other things. I'm not going to go through this in much detail, but this type of optimization, if you get proxies where you can do the 10,000 iterations, will give you a solution that minimizes the impact of model error. This is really what you want to do because you have so many moving pieces. You want to come up with a solution that's robust with respect to model error. I think I've coined the phrase, SEME, solution elasticity with respect to model error. You arrive at a solution that's robust against model error and then you know, in some sense, that even if all your models are wrong, and you've led them all five standard deviations, this is the thing that's going to risk manage you the best. At Stanford, we spend most of our time billing proxies from software on the asset/liability side and trying to find dynamic strategies. You get back to the Tenny Gerard question of how do you look up all strategies. I don't really think you can look at all strategies. You can optimize broad classes of dynamic strategies using this through what is called the genetic algorithm, revolutionary strategy operators. When you get to the end, and when you have something that you can do a lot of computations on, you can find a solution.

What does this look like? Once you get it down to this level of compactness, the curve is really small. It can be a java aplet. The average size is about 47 K. That's a small piece of code. The way we do it is in a genetic algorithm, you take the trial solution so there's a trial solution. Now if we're in a trial solution, you generate 50 or so random, stochastic interest rate paths. For each path, you then loop through time. At each point in time, you compute interest, principal, value, and liabilities, and you apply the strategy. You buy, sell, and keep doing that. When you get to the end of that scenario, you compute summary statistics and store them in an array. You go back and do some more scenarios. After you have reached the end of that array, you assess what your goal variables have done across all the scenarios, and then you take that back and you find out what your evolutionary strategy or your genetic algorithm says about modifying the solution to get you closer to your objective. You do this 10,000 times, and it takes four-and-a-half minutes to run. It takes cash-flow testing from something that's maybe done once a month, once a week, or once a year, and it converts it into a tool for the real time investment decision maker. Once you get a proxy from the asset side, and once you get a proxy from the liability side, then your investment manager can balance a strategy, a change in strategy or an available set of assets against that and find out what it does, if anything, to the solution and what the elasticity of that is with respect to model error. The goal from our point of view is to take the cash-flow testing process and move it fully to a dynamic adequacy analysis where you can do a computation very quickly. I mean even if it takes a couple of hours overnight, the market is moving faster than that.

I'm going to talk a little bit about yield curve rotations, which Teri also mentioned that. Basically, in this process that I just described, you have to generate an interest rate scenario. All the modeling you do is conditioned on the yield curve shape. All of your lapse rate assumptions are condition responses of some sort to the shape of the interest term structure. You can combine economic variables. I suppose that could be expanded, but in general, my experience has been that you'll be doing a good job if you can incorporate shifts in rotations and the yield curve.

Why does the yield curve change through business cycles? I think it's important to think about this, because if you're going to generate scenarios, you could actually defeat the whole process by having a large number of the scenarios be things that just couldn't happen. When you generate interest rate scenarios and you look for model error, you want to generate these in a position that's consistent with history. There has been many techniques to do this. There has been many time-sharing studies, and many garch techniques. What kind of numbers are we dealing with? If you model basically looking at serial correlation between the long end, you'll find out that you end up with an R squared, which is 12 or 24 down there? If you take a look at the correlation between the long end and the short end moves since 1979, you end up with about 224. So the point here is really that whatever you're going to do, it's important that you do something more than just parallel shifts because these things are not always moving together. Now I'm not going to talk in as much detail, but I'm going to talk a little bit about the yield curve and create these proxies and how you want to model things. The idea is to take all your classes of assets and all your liabilities and somehow reduce them into a manageable group and have a compact, efficient model of that manageable group that replicates very closely what you would have gotten had you run an extensive run on, for example, Bloomberg on the liability side. Teri talked in her presentation about the types of risk. You have default risk, which is credit risk, and you have interest rate risk, which is a sensitivity with respect to interest rates. Of course, the more random thing you have is something that I call liquidity risk, which means that maybe that product has fallen in disfavor in the market and things have widened out. Whatever it is, default is certainly a primary example.

How would you want to make your modeling a default risk across your fixed-income securities be consistent? From 1979 to 1999, you'll see that you can take a total return cumulative growth on the Lehman intermediate corporates. That is almost exactly replicated by a portfolio. It's 20% for the three-month and 80% for the long bond. Now that's significant because whenever you generate these scenarios and develop model error for corporate spreads, you want it to be tied to this empirical relationship. So that's another thing. You have compactness and you also want to capture empirical regularities to the extent you can. You can do the same thing by taking a look at a mortgage index. I know people get up every day and look at the mortgage sector, the corporate sector, and the agencies. They say, I wish I had a nickel for every time I picked up something from the street that said, mortgages cheap to corporates. Then we see the graphs and charts and the cheapest is 15 basis points. The 15 basis points are way beyond any model error you're dealing with, which is so minuscule that it doesn't matter. I mean if you're devoting your resources to figuring out whether or not corporates are 15 basis points, I think it's misguided. You should be sure that when you model across horizons, you get a good relationship. It turns out that the mortgages again can be pretty much approximated by the three-month and the long-term which is the short and long on your curve and it's the 15% and the three month, 85%.

There are variances as you go through time. They sort of revert to each other. So if we have a block of corporates in our portfolio, we'd like to come up with two or three compact representations, depending on maturity, where we're dealing with model error in that magnitude. So the question is, how much should I randomly vary that spread when I'm doing this testing. How is it related to my yield curve?

I said a cash-flow matched portfolio is duration matched, but the other isn't true. What makes all of this important is not necessarily that you were in a Treasury portfolio for 20 years. You would pretty closely track a comparable proxy and the corporate. What drives the problem is that at any given point in time, if you're in something that looks like an index, you might have periods of time where it's all under water. You might have periods of time where it's looking pretty good. You might have periods of time where it's just sort of mediocre, but the purpose of asset adequacy is to minimize this residual risk and this mismatch between your assets and your liabilities. By using a tool like this, you really score points in the long term by understanding model error and understanding your asset/liability match or mismatch, and not by doing relative value analysis on a daily basis.

I'm not going to go back to some of the other slides we had. In summary, I'd like to just give you this macro picture. You have things that generate asset detail and liability detail. You reduce those down to econometrics proxies of things that replicate them well but have model error. You model the error and you look for dynamic strategies as you go through time. That gives you something that you can do with model error. Here's where the cleverness and the innovation comes in. You must find that strategy that is robust with respect to model error so that the decision you make is good within what you foresee to be your range of model error.

MR. JACKSON: I have a question for Ted and it directly relates to use of this kind of tool. Who uses it and when is it used? I guess I'm thinking in terms of a paradigm where the investment and portfolio manager in our firm has cash to put to use. I'm not quite clear as to whether your process would be used annually or whether it's used each time there's time to purchase a security. How does it actually work?

MR. CACKOWSKI: The way this is used is you take whatever the last asset run was for BondEdge. You take a look at what the last liability run was from whenever they do that. They might do it weekly, or they might do it monthly. You create the proxy. The way we envision this is you take this proxy and you post it to a Web server on the Internet, and then you have a java aplet where you put some modeling information about things that are available. The java aplet does the 10,000 iterations where it determines whether or not the things that are available, given your dollar to be invested, competes well against your profile and what your model proxy looks like. The answer is, it's kind of contemplated to be a real-time decision updated as often as you're willing to update the proxies. Then make investment decisions based on those proxies. So when the java aplet comes down, it takes about a minute and a half. You can click and shoot security that might be available, and find out whether or not it ought to displace anything or whether it ought to be an addition to your portfolio for new money. So real time is the answer, Rick. We can talk about garch for a few minutes.

MR. JACKSON: I do have one question for Teri as well. You did mention this, I see our portfolio managers in my firm being focused on duration targets. You threw out a concept that was a little bit innovative, and it sounded like you were talking about your uncertainty measure for prepayment. Do you have a durational uncertainty measure for specifically mortgage-backed securities? I see us trying to target duration in our firm with some of our clients. With mortgage-backed securities, it might make sense to consider a duration uncertainty measure.

MS. GESKE: We actually do that. The prepayment uncertainty measure that I described looks at the sensitivity of price given a change in your prepayment model, but you can do very similar exercises to derive the change in your duration given a change in your prepayment model forecast. We've done some work on that. I actually have written a chapter that's going to be appearing in a book called *The Handout of Financial Risk Management*, put out by the Global Association of Risk Professionals. I show how your duration can change given a 10% change in your prepayment model forecast. It's a related way of looking at it. I think it's a good way of thinking about duration as being something that's within a range instead of a point precision measure.

FROM THE FLOOR: Is it fair to say that duration is more or less sensitive to model error than price?

MS. GERSKE: Is duration more sensitive to model error than price? I would have to think about that. You can have instruments like interest only (IO) securities or principal only (PO) securities, where the durations are extremely large, and their convexities, which is another way of saying the instability of their duration, is also very large. In those cases, I would suspect that the duration sensitivity is probably larger than the price sensitivity, but I'm not sure that would hold in all cases. I have to think about it.

MR. CACKOWSKI: You have two inputs there that are going to affect the price, and the price is an intermediate step in getting the duration. As the prepayment assumption changes, the average life changes on the security and as the average life changes, the pricing spread is also going to be sensitive to that. The answer is they can very well compound each other; it could hit you both ways.

MR. JACKSON: I'm going to throw out one more question for Teri, and then if we don't have any further questions, we'll close it out. Teri mentioned the volatility of credit spreads and how corporate spreads have historically been there. I guess one question I do have for you that I'm curious about is whether you've observed a constriction or reduction in corporate spreads in total. I guess I would be focused more upon some mean results. Some of the pricing work that's being done for annuities and other products kind of requires a certain spread to be earned. Have you seen or observed any reduction in spreads over time as interest rates have come down, or is there kind of mean reversion going on here that comes back to averages?

MS. GESKE: I've been involved in this area for seven years and for six of the seven years, there was very little volatility in credit spreads and a general tend toward narrowing. Ben can probably back me up on this. There was on the order of 10-15 basis points in spread volatility over the course of a year for most types of credits. Last year we had a 50-basis-point widening in the course of a week. Although it has come in somewhat from that, it hasn't reverted back to some long-term mean yet. We really don't know if we've entered into some new era of wider credit spreads or if we're entering a new era of increased volatility. With respect to credit spreads, the period of time over which you would measure averages, standard deviations and mean reversions is going to be incredibly sensitive to whether or not you include the fall of 1998. I would say that for six years, it was very stable, and the flight trends were narrowing. All of a sudden, it just blew out and has come in somewhat since then, but not by very much. It is a very unusual situation. I'm not sure what conclusions we can draw just yet.

MR. CACKOWSKI: One of the problems I have with value risk is it's largely dependent on the stability of the correlation matrix and, of course, that's the risk matrix popularized by J.P. Morgan. You take those things and you extrapolate that out over a year, two years, five years, and ten years. I think you'll end up with something that I'm going to call noise. So if you have enough model error and your paths are random noise, you end up with noise. One of the things that I believe you need is to get your interest rate paths concise, and, in some sense, consistent with history.

Then you need to take a look at empirical regularities. We talked about the corporate spreads. As you notice, if you own 20% in the three months and 80% in the long bond, you're going to have a portfolio that's going to keep crossing something that's just basically representative of Lehman intermediates. The answer is, I think you're right.

MS. GESKE: I'd just add that the value-at-risk was derived to be used in a bank environment on a trading debt where a horizon is a day or two days or a week at most because volatility with respect to your trading portfolio could put your bank under in the course of a few days. You have to look at what your horizon is. Your horizon probably isn't three days so you have to adjust your value-at-risk inputs to reflect your horizon. Then ask yourself whether or not VAR is the appropriate tool for the horizon that you're talking about.