RECORD, Volume 22, No. 1*

Marco Island Spring Meeting
May 29–31, 1996

Session 84TS
Algorithmic Aspects of Interest Rate Generators

Track: Investment/Computer Science
Key words: Forecasting, Valuation Actuary

Moderator: MICHAEL F. DAVLIN
Instructors: SARAH L. M. CHRISTIANSEN
MARK TENNEY†
Recorder: MICHAEL F. DAVLIN

Summary: Panelists share their knowledge of the following in this teaching session.

- *quasi Monte Carlo random numbers,*
- *log normal and other probability distributions of interest rates,*
- *various generally used algorithms, examples, and applications, and*
- *strengths and weaknesses of the various algorithms.*

Mr. Michael F. Davlin: Sarah Christiansen is the assistant corporate actuary with the Principal Financial Group, and she recently wrote a paper that is going to be published in *The Random Character of Interest Rates* by Probus in Seattle, Washington called the "Representative Interest Rate Paths." She won the Actuarial Education and Research Fund (AERF) 1996 award for that paper. She is going to talk about selecting interest rate scenarios once you've generated them. Mark Tenney is going to talk about some of the more algorithmic aspects and some other approaches to selecting interest rate paths from sets of interest rates.

We're trying to find a way to avoid running thousands and thousands of scenarios in order to get a reasonably accurate result.

Ms. Sarah L. M. Christiansen: This session was originally planned for people with no experience with interest rate scenarios at all, so I thought we'd try to start from

†Mr. Tenney, not a member of the sponsoring organizations, is President of Mathematical Finance Company in Alexandria, VA.

the beginning. We will get to the represented scenarios, but I want to start with just the basics.

What is an interest rate scenario? Basically, it's a set of interest rate curves. One curve exists per period in your time horizon, whether your period is a month, a quarter, or a year. Your time horizon usually goes out 20 or 30 years. I'm going to go a little bit further and define a curve with 12 maturities ranging from 3 months to 30 years.

Why do we have interest rate scenarios? We have a number of people who are interested in them. We need to have confidence in and an ability to handle the future; as such they are the foundation for many models, some of which are wanted by regulators. Regulatory agencies are interested in the results of cash-flow testing, and management can and should be interested in the results of cash-flow testing. Whether they actually are interested or not may depend on your management.

The whole idea of an interest rate scenario is that the curve will change levels, and it may also change shapes. Cash-flow testing is the big reason that actuaries originally became involved with interest rate scenarios. It started a little over ten years ago with New York Regulation 126. At that time, the New York insurance department was very concerned because we, the industry, were selling annuities with guaranteed rates of 10% or more. Interest rates were high, but the assumption was we would be able to invest forever at 12%. Regulators thought that didn't make sense. Originally, they gave a choice. Statutory rates were high, so they gave companies a choice between doing cash-flow testing on at least seven scenarios (which they specified and I'll get back to those later) or using a lower rate for reserves. Now, there’s no more choice. If you're going to do business in New York, you're going to test on at least the New York seven scenarios.

What are the New York seven scenarios? First, they're all parallel shape. They never change the shape of the curve. With the first one, there’s no change at all. Then there’s one that starts and rises 0.5% a year for ten years so it's 5% higher than the original curve at that time and stays level forever, or until the end of your time horizon. Then there’s the opposite scenario: it drops 0.5% per year for ten years and then remains level. Yet another goes up 1% a year for five years, back down 1% a year for five years and returns to the original level. Another is just the opposite: the rate drops first for five years at 1% a year and then comes back up at 1% a year for five years and then stays level. There are two more. There’s a pop up, where the first year you go up 3%, and there’s a pop down, so you go down 3% the first year, and then you stay level at the new rates forever.

New York also said every scenario has to have the current curve for time zero and that's kind of important. So that was the start. They are not saying by any means that those are adequate or those are all the ones you should test, but they are saying that those seven scenarios should be included in your cash-flow testing.

Now the Standard Valuation Law from the NAIC followed on the heels of New York. The NAIC came up with a Section 8 asset adequacy analysis. This does not always require cash-flow testing, but it is perhaps the most common method of satisfying it except for some small blocks of business. There are other ways around it, but cash-flow testing is a big area. There are two Actuarial Standards of Practice, numbers 7 and 14, that relate to cash-flow testing.

In addition, you can do some product pricing. Our company uses cash-flow testing for product pricing to test whether they're really going to have the profitability they think they're going to have or would like to have. You can use it for portfolio management, and you can test a reinvestment strategy. Which reinvestment strategy works better? This one or that one? The prediction is that rates are going to rise, but we're not going to make bets on interest rates. I don't know how many of you read the article in the May/June 1996 issue of *Contingencies*, in which there was a review of *Innumeracy*. It talked about a scam. A company could obtain a list of wealthy individual investors and send them a prediction on the Standard & Poor's (S&P 500). Let's say that there were 128,000 people on this list. They could take 64,000 and say rates are going up or the S&P is going up. They could take the other 64,000 and say rates are going down in the next week. So they could say not only what rates are going to do, but when it's going to happen. Then, the next week, they might take only the list of people that they had made a correct prediction to, and divide that in half and send half an up prediction and half a down prediction. This could continue, and after about four or five weeks they might say, We can no longer continue to give you this service for free. We've given you correct predictions for four or five weeks in a row, and we would really like $1,000 from you for predictions for the rest of the year." The investors might figure they'd received correct predictions for five weeks in a row—that's a pretty good track record! Did they know about the wrong ones? No. So even if we do have a fortune teller I don't think that's what we want to rely on.

We have many choices for interest rate generators. What way do we want to go? We have a choice between deterministic (the New York seven are definitely deterministic) or stochastic. If we're going to use stochastic, do we want to use arbitrage free or binomial lattice or some continuous models? Arbitrage free can be either binomial or continuous. What is arbitrage free? Arbitrage means that the expectations theory of the term structure of interest rates is met in your scenario. So there are some implied forward rates. Your scenarios average out so that time equals time—that's the fundamental equation in arbitrage-free scenarios. Perhaps as you look at our scenarios, you see that they were years in the future. As you look at the implied forward rates and other things, the rates implied for a year from now are those that really come to pass on the average over the scenarios. That's the sort of thing that an arbitrage-free set of scenarios requires. It's very much tied to what today's curve is. That is very appropriate for selling derivatives. It's very appropriate for finance people and Wall Street people. They live and die for arbitrage free. If you don't know that you should have arbitrage free for pricing these things, you could be taken advantage of. They don't agree on prices, and that's another story. On the other hand, there is a big question in the actuarial circles: is this appropriate for valuation purposes? Is that assumption appropriate for doing adequacy analyses? I don't think so. There are many people who don't think that's an appropriate assumption for that purpose.

That brings us to the next question which is, why do you want to do it? Are you looking at what price an asset or liability will bring? Well, if you do the contingent claims analysis and you buy all the rest of the assumptions like liquid markets and transaction free and no tax constraints, then maybe that's right. Also, what are the types of rates? Do you want yield-to-maturity rates? Do you want spot rates? Do you want forward rates? Yield-to-maturity rates are the type of rates you see on a home mortgage. The same rate is used for the entire contract. It's the rate that your Treasury curve is expressing when they say that the 30-year Treasury is 6.94%. It's a yield-to-maturity rate.

Spot rates are the rates for the zero coupon bonds. Sometimes they're called pure discount rates. You don't receive any coupons and then everything is paid back at maturity. The one year forward is the rate you might expect a year from now to get you from time t to time t plus one. There are implied forward rates and there's a one-to-one relationship between spot and forward rates. Going between either of those and yield-to-maturity depends on the cash-flow pattern you assume.

We're still bombarded with choices. Do you want to use a discrete model? A continuous one that terminates? What about rate by asset type? Do you want to have different rates for corporate bonds, commercial mortgages and your residential portfolio? And what about Treasury rates? My investment people tell me that if I'm valuing my portfolio, I really want to look at each of those rates for cash-flow purposes. First, I want to look at assets and liabilities on the same set of curves, and second I want to look at what I'm going to be getting for reinvestment. What's a realistic rate if I'm not buying just Treasuries? Are Treasuries the right choice? We tend to blend our rates based on the proportion of each of these types of assets. When we do the liabilities, we use a blended rate. Spreads. I don't believe that spreads over Treasuries are constant for the various asset types. I don't care what commercially available projection software does. They don't stay the same. They're not going to remain constant over time, and they tend to vary with maturity as well.

What are the kinds of models? There are discrete models, which could be binomial lattices. They're often one factor. If it's a one-factor model, it's usually a one-month forward or a one-quarter forward. They're connected. By that we mean that up and down, there are two choices at each point in time. Interest rates go up or go down. An up/down will get you to the same place as a down/up. You have 2^n possible paths and you'll be selecting from that. That's why they tend to look at paths. That's where the word path comes from. Academics really like binomial lattices. I won't say that nobody in the industry uses them because there are industry people who do. The Government Actuary's Department is one but, for the most part, academics tend to prefer binomial lattices.

For the most part, the industry tends to prefer continuous models. One normal one is supposed to be the continuous analog of the binomial lattices, but if there’s a possibility that your rate won’t move by a single jump, it actually could stay the same. It could go in between. Log normal often will have a mean reversion. You might have a multifactor model with at least two rates, and you might use correlations between the rates as you run your model.

Let’s discuss a sample binomial lattice. The big thing is that this is connected. If you have end time periods, you will have end plus one node and 2^n different paths. You can see something that's going to happen. Even if you assume this is multiplicative, so that when you divide by one plus the rate it will never become negative, it's going to get very close and tight toward the bottom. It doesn't seem to have any limit on the top. We need an injected dose of reality as we are going through our generators and looking at them. We want to make sure that we don't get negative or excessive rates.

I tend to think that there's politically a good reason to assume that there's going to be a floor on rates. We have inflation and big debt in the U.S. We're never going to see 0% interest. By the same token, I don't think the U.S. would be happy with rates over 25%. I think that whoever is elected is going to get dumped out and whoever is Federal Reserve board chairperson is going to get kicked out. We do have some political and economic pressures that say the rate will remain in balance. They might push the balance, but they're going to remain in some balance. We need to look at some realistic possibilities for what the curve shapes are. I think the shape is something that really has not gotten much attention. Also consider computer time. Whatever you're going to do, you need to look at the time it takes to develop and maintain your generator.

Rates for the key maturity point is modeled first. You're going to have a volatility parameter that we'll come back to in a little bit, and that sort of says how much movement can occur during one time period. Basically there's a standard deviation of the movement, which is going to show you how much randomness there is in things. What are the common choices for rates that are modeled first? Short rates are very common. There are also long rates. Some people even model inflation for some real rate of return. We mentioned a log normal generator. The log normal depends on the assumption that you are working with a standardized distribution:

$$R_{t+1} = R_t \exp(\sigma Z_t)$$

And you could do some backward substitution and find

$$R_{t+1} = R_0 \exp(\sigma \sum Z_t)$$

Well, even if the *Z*s are reasonable, there are some that can get very large or very small. Second, they're kind of random. Ideally they would be random, but we need to be able to reproduce the scenarios and no computer or any random number generator buried in the computer is actually producing truly random rates.

I work in APL and we have a random number seed that actually changes every time you make a call to the random number generator. We actually set it up with a random seed for starting based on the time and date of the run. We keep track of it so that we know what it is; that way if we have to reproduce the scenarios, we can. Well, the problems with the log normal are that those x points get to be a little bit too high or too low and give you rates that are too high or too low. The first thing that one might look at is putting a barrier on. You really don't want to fall off that thing and let the sharks eat you. If you use an absorbing barrier, you get very sticky stuff because you'll get the same interest rate time after time after time. Or they can be reflective and bounce back. Mean reversion is a solution to this. It pulls toward a pre-set goal. Basically you're going to have to set the goal one way or another.

There are three possibilities that make some sense for setting that goal. The first one is to use a starting rate. Given the source of interest rates, a textbook will tell you that you don't have any better guess than what things currently are, so you may as well use the current value. If you were trying to do an arbitrage-free generator, there is a rate implied in the current curve for the maturity that you have selected as the seed. That's a second choice for what you should use. Maybe you have an expert who knows where this rate should go. Perhaps you want to use that expert opinion. You choose. It's nice to have it stay the same. They work by disrupting the addition of exponents. If you disrupt the addition of exponents, then you don't get these arbitrarily high positives and low negatives that will give you very high or very low rates. This is often done in addition to the boundaries.

Suppose you're looking at a two-factor model. What choice might you want to make for the second factor? Mr. Jetton used the deterministic calculation with volatility of a secondary rate. Mr. Strommen in his response said, well, I use a ratio. Mr. Gurski used correlation coefficients, he was the only one who used a set of three intermediate rates, and he puts some spread constraints on. He said that if the difference between the ten year rate and the one year is over such and such, I'll reduce it by 50% of the excess. And the same way with the 10 year and the 30 year and he had some correlation coefficients between them, and Mr. Mereu did inflation first and then picked out some other rates. How do we do the rest of the yield curve? Well, we have some published generators, Mr. Jetton and Strommen had weights, Mr. Mereu had a formula, Mr. Gurski fits blinds, and I use a shape with a Markov chain process and this is nothing more than a random walk.

I'm going to give you a little bit more on my generator. I like to key off the 30-year maturity. Volatility tends to be greatest at the short end of the curve, and smallest at the long end of the curve. I like the least volatility. I usually use a 30-year maturity; however, it is possible to reset mine to use any of the maturities. Things could bounce out of an expected range and your volatilities won’t follow as nice a pattern if you use something other than the longest rate. I use a lognormal with mean reversion and run on random numbers. In studying Treasury rates, I took the historical Treasury and looked at the curve as ratios. In other words, I divided everything by the longest rate so that I obtained numbers that were in the range of 0.4–1.4. As I looked at these Treasuries, I identified half a dozen shapes. Actually, there are more than half a dozen. I started out with seven, and I'm now at eleven. They range from a very steep normal curve to a very inverted one to one that is flat. It’s what you would call a bump and a mirror image, and what I would call an early gully. They're not quite level, but it looks like maybe a sine curve or a cosine curve without too much wigging and you might start out up or you might start out down. I numbered these, 1 through 11.

I listed the original curve, I divided all the rates by the long rate, and I tried to determine in which one of these shapes this fit best. Then I have a random walk matrix that's a little bit mean reversionary. We have a little bit of a push down toward the normal, upward sloping, and we do a random walk and shape code. You pick a uniform random number. You have your curved shape, and you look at a cumulative probability distribution that gives me this shape. I have this shape matrix. People say I have an 11-point or a 2-factor model. One factor is a shape code. The other factor is a long rate. I specify the rest of the points tentatively as the product of the long rate with the factors for that shape, so I instantly have all 12 points on the curve. I didn't find that the Treasuries matched these shapes perfectly all the time, and I don't want mine to be perfect shapes all the time either, so there's some necessary adjustment. I'm going to limit the change from year to year by a maximum percent of what I was at before. By the way, that maximum percentage depends on the maturity, so it's bigger at the short end and smaller at the long end.

I also set some boundary conditions. New York originally had a 4% minimum on theirs and then they decided that wasn't low enough, but basically it's at about half the original curve or 3%, whichever happens to be less is my minimum and a flat maximum of 25%. Not all the shapes are perfect. The shape proportions are basically comparable to historical, at least as well as I can test. I looked at the percentage of inverted curves and compared it with David Becker's work at Lincoln National. I generate spot curves, but I don't get negative forward rates and I can convert to yield to maturity.

Now suppose you have a generator and you want to do representative scenarios. We'd love to do lots of scenarios but we have limited time and resources as Mike alluded to. In the real world, I'd like to do a thousand, but not surprisingly, time will only let me run 50. I'm going to do the New York seven, and my two interest rate shocks will be up or down 1% tomorrow. I'd like to reduce the number of those scenarios. I'll accomplish this by finding a representative subset. How many subsets are there? This formula will tell us:

$$\left(\frac{1000}{50} \right) = 9.640461 \ x \ 10^{84}$$

I really don't think I want to look at each possible subset. That's considerably worse and not better than looking at my thousand scenarios.

What do I mean by representative subsets? For each maturity, the subset has approximately the same mean, range, and variance. The subset is going to have exactly the same curves. I'm going to set up for 1,000 reproducible scenarios, and keep track of my random number seed. I'll run 200 at a time and pick 10 representative ones.

For each maturity rate, I have the maturity rate for each scenario for each projection year. I'm going to start with a three-month rate and I'm going to find the simple descriptive statistic for that three-month rate. For each scenario and for the set of all 200 scenarios, I'm not going to consider every possible subset. I still have 10^{16} subsets. I'm going to create a candidate list instead. The candidates are going to be subsets of scenario numbers. I'm going to match extremes. By that I mean the minimum and the scenario that I pick is going to be the overall minimum for that three-month rate for that set of 200. So if the minimum is found in scenario two, six, and ten and the maximum is in five and 141, I'll look at two-five, six-five, ten-five, and two-141, six-141, and ten-141. If, by some chance, the same scenario happened to have both the minimum and maximum, I will not put it in twice. And that was something that somebody pointed out to me and it was a good catch, because we actually found out we did have that. What we're looking at is the minimum of minimums, basically, and the maximum is the maximum of the maximums.

Now I'm going to look at the first scenario. I have an overall average, and an overall standard deviation. So I can calculate these four values, µ plus 0.85 sigma and minus 0.85 sigma and µ plus and minus 0.65 sigma. I'm going to look at my list and find out which scenario that I don't already have in here is closest to each of those values. I want to add those numbers to each one of these subsets. Now each subset has six elements and a mean.

I'm going to pick four more scenarios such that I'm forcing the mean of the ten to be as close as possible to (10µ - 6m)/4, where m is the mean of the 6 elements previously chosen. So I'm going to look at those and add them in and those give me the candidate list for the three months.

I'm going to repeat that with the six month, the one year, the two year, etc., and that will complete my candidate list. I'm only going to look at those candidates. I must have at least 12. I must have at least one for each maturity. It has been my experience that our candidate lists run from 100 to 200. That's a far cry from 10^{16}. The goal is to match simultaneously the run statistics with those of the subset for all maturities This is very important. I'm going to weigh those maturities arbitrarily depending on what I'm doing. If I'm pricing, I may say that the weight in the three, five, and seven area may be most important and others might be less. Weigh them from one to four arbitrarily. Determine the statistics for that and for the run. I'm going to create a choice function which is nothing more than a least squares difference weighted by my weight. I'm going to choose the value where that choice function is the minimum. The choice of the weights is sensitive, and you can add some things into that choice function if you need a match in extreme. If you want to match medium, add them in. You can add things, but you'll just have a bigger number to choose from.

There are repeats of the other runs. Keep track of the scenario numbers that will be run from the winning candidate each time. Recreate the 50 scenarios and compare the descriptive statistics. It worked out very well for us most of the time. We get an automatic comparison produced by the generator. You can't see the data, but we do have the original data, or the representative data.

In June 1995, we had no difference at all; we hit the minimum everywhere. The mean line is very close. The median line drops, and the one that drops has a minimum of about five years. For the maximums, it's flat for the first ones and then bounces up. The standard deviation is on top. I expected that the representative would always have a larger standard deviation because there were fewer points than the original. Indeed the data were exactly the same; there were less data. So we did very well and we found a new algorithm. We're reproducing, but we're not determining probability of each scenario; rather we're reproducing the probability distribution of the scenarios of a set of 1,000 and a subset of 50. We actually met the target.

From The Floor: You looked for 100 to 200 scenarios. Do you have some sense as to what degree the cumulative distribution function was covered?

Ms. Christiansen: I don't have a real good sense; Steve Craighead of Nationwide had that opportunity. He gave this a workout that I was unable to do. He runs 10,000 interest rate scenarios on 31 lines of business. He did more than this because I had tested this on not just interest rate scenarios. We get the same results from our cash-flow testing on these interest rate scenarios once you've cut it back. We were doing fairly well on that. He did this and ran through his cash-flow testing on all 10,000. He cut his 10,000 down to 400 using the small set by grouping in 250s and taking 10 out of 250. He said it worked very well in reproducing the entire spectrum. He did find that he did need to adjust his weights. I do recommend that your weights depend on what you think is important for what you're doing. For example, even though the 30-year rate is the key rate when I generate the 30-year rate, if I don't have liabilities out there, if I don't have many assets out there, that should be rated a one for least important. If your assets and liabilities are concentrated in the three- to ten-year range, maybe you ought to give those weights of four and look at the one and twos as maybe a three and your 20 might be a two, and your three-month rate, that's not going to make a whole lot of difference in anything. So you might look at things like that and start there, but things are very sensitive. I'm more concerned about matching the minimum service first ten years, so you might want to put some sort of constraint on there.

From The Floor: The selection of your scenarios all boils down to what inputs you're putting into a lognormal generator. Maybe you can give some brief comments on choosing the volatility factor. I know some of the software companies recommend a historical range of values that might approximate what that volatility factor is. I want to get your comments on how to choose that not only today, but maybe 20 years from now, based on the past histories.

Ms. Christiansen: I have worked with some of the people on option pricing, and we kind of looked at some of the historical rates. We actually have someone working on that right now. Volatility factors do tend to decrease, and tend to be inverted relative to the level to the maturity. I put in higher volatilities for the short end. If I were using a one-year rate to key off of, I'd be using something much higher, maybe 14, 15% on the 30 to maybe 25–30% on the three-month rate. I tended to put in something that was linear and basically it's only used for the key rate so once I picked my key rate, I just pick out the associated volatility factor. The rest of the volatility comes from the shapes.

Michael L. Yanacheak: Of the two factors that you modeled, the shape of the curve and the 30, was there any correlation between the shape, given that you already calculated a 30-year rate at all?

Ms. Christiansen: I didn't put any in, but it does seem to work with other models and I don't know why.

Mr. Davlin: Our next speaker is Mark Tenney.

Mr. Mark Tenney: I'm going to talk about construction, calibration, and essentially implementation of interest generators. An interest rate generator generates random sequences of yield curves, starting with the current yield curve. If we're dealing with an arbitrage-free model, we see the arbitrage-free model has several components, one of which is the possible sequence of yield curves, or more generally the possible sequence of events or market prices. We have two sets of probabilities for the same set of events. One set of probabilities is the realistic probabilities and one set of probabilities is the risk adjusted.

So you can think in terms of the prototypical arbitrage-free model, the option pricing model. Within the risk-adjusted probabilities, the expected return on the stock is the risk-free rate. Using the realistic probabilities, the expected return on the stock is greater than the risk-free rate. So when we build one of these interest rate generators, we want to cover both of these two because they have applications to different things. In one case, we're doing prices, one you use the risk-adjusted probabilities like Black-Scholes did. You would expect that the return equals a short key rate. If our application is portfolio analysis, optimal portfolio, or whatever, we want to use the realistic probabilities like Markowitz who analyzes portfolio choice. I guess the other points here are that there's the issue of calibration to an initial yield curve. If you have an initial yield curve, you can calibrate either of the variables of your model, keeping your parameters constant, in that case you're not going to exactly fit the curve unless you picked in some extra parameters to do that. As you run through your scenarios, only your variables are going to be changing since your parameters are fixed. If we're thinking about some natural historical data, then you'll match an initial curve with an arbitrary set of parameters. That's going to give you worse performance from matching further curves later on holding those parameters constant, which is what we're essentially in the situation of doing when we're using these scenarios given today's information.

We have these basic applications which I just covered. For asset/liability management, we can look at something Sarah has talked about, but we can represent our results either in terms of average and standard deviation or something like distributable earnings. We also can concentrate on the tail end of the distribution, the probability of ruin. If we're doing product pricing (and I'm talking about the insurance product,) there are two ways to do it. We can either use scenario analysis, cash-flow testing, and use the realistic scenarios with realistic probabilities, or we can price the liability just as though it were an asset. We used the risk-adjusted scenarios.

If we're looking at this from an institutional point of view, what do we have to do to develop an interest rate generator? It's a little bit more than just coding up a random number generator and having some additional multiplications and divisions. First, we have to determine what we're going to use it for. We have to decide if we're going to have one common generator for the two different applications—pricing and cash-flow testing—or if we are going to have two different generators. If we use different generators, we start running into consistency problems. The best solution is to have one arbitrage-free model, and use the two different probabilities: realistic probabilities, and risk-adjusted probabilities. Once we've made that decision we go on to build the core interest rate generator itself. That's what is going to generate our variables. From those we have to determine the yield curve and we have a pricing system that also priced some other securities, a bond, or whatever.

In order to run this, each day we must take the Treasury curve, and we take our parameters and set our variables to give the best fit to that curve. If we're a little bit more loose about our procedure, then we go ahead and fit many parameters in order to fit that curve as well, which is the approach that Wall Street typically takes. Even though we're going to do it right, we still have to do a historical calibration to figure out what those parameters are. We do a historical study over a 20- or 30-year period, and find some parameters that allow both our probability distributions to work. We want to be able to match the Treasury curve relatively accurately, say month by month over the last 25 years. We also want to get the probability distribution of short-term rates, yield-curve inversions, and long-term rates to be relatively accurate. For runs testing overall probability of ruin and asset/liability strategies, we would be using the realistic probability, and for product pricing we could use either realistic or risk adjusted. The final step is to develop a methodology in which we can use fewer scenarios in order to get the results we're interested in. Once we've done that, we can go ahead and try to add some additional variables to our model. There's inflation, an equity-index, and a dividend yield that would allow us to do things like equity-index annuities, which depend on both your interest rate model and your stock prices. So the components of our complete system are the core interest generator, which is generating our yield curves, our daily calibration software to fit to the Treasury curve, and maybe we have some other things we want to fit. There's also our historical calibration software.

Let's discuss the components for the core generator. As I said before, we have some core variables that change randomly. We essentially have this stochastic process in risk-adjusted form and risk-unadjusted form, which gives us the two different sets of probabilities. Given our model, we then get a bond pricing model so that we can calculate yield. We can add our derivative model on top of that. So the software we need must perform the stochastic generation of the core state variables and calculation of the yield curves for each set of state variables at each date. When interest rate generators are discussed they often think this is the whole story. As we've seen, there's actually much more that takes much more time. We do our daily calibration as we change the yield curve from day to day; our variables change, but our parameters are fixed. Many people have taken one-factor models that can't possibly describe yield curves over a ten-year period or even a five year period. They just recalibrate their parameters each day and ignore the fact that their parameters are changing from day to day.

Maybe we're doing cash-flow testing for a ten-year horizon or we're thinking about pricing a derivative. We plan to hold it longer than the time it takes us to sell it to the next customer who comes along, or to lay off our positions. Then we need to think about a longer term horizon and we need a model that's stable over that longer term. We're particularly worried if we do our cash-flow testing this year and we decide our institutional strategy is to hold these assets and these liabilities. If we're using one of these models that has these fudge parameters, then next year we may have basically said that last year we had a terrible strategy. We need to change it, but next year it ends up being the same story. If our model isn't time invariant, if it isn't stable, then the decisions we make as to what to do are not going to be stable. We get into a situation of spending three months doing cash-flow testing, coming up with one set of decisions and then essentially changing it the next year; however we may have lost money in the meantime.

Let's go into some depth on the historical calibrations. We must determine the Statistical Criteria to use for historical evaluations like David Becker's Quantitative Stylized facts. We look at the frequency of yield curve inversions over the last 40 years which is 15%. We take the three-month yield minus the ten-year yield.

Becker also looked at how often a ten-year yield is between 10% and 12% or how often the three-month yield is between some different range.

Sarah has already talked about scenario reduction methods. Let me just make one comment. When many people talk about scenario reduction methods, they are really talking about methods to reduce the number of scenarios in order to calculate a single number like the price of a security. So you can use an antithetic method or whatever. Ideally, you would be able to calculate the quantity you wanted with two scenarios with zero variance. From the point of view of cash-flow testing, that's not what you want. The conclusion is not that your portfolio has zero variance. What you want to do is find out what the probability distribution of your portfolio return is, so it's a much harder job. Sarah's approach is one of very few approaches that address that question.

So what are our requirements for a stochastic interest rate generator? First, we want to include all the yield curve sequences that could happen. We don't want the impossible ones. We want to have the correct probability of a sequence occurring for our realistic probability. For our risk-adjusted probabilities, we get a different probability—one that incorporates, in essence, the market price to risk. We want to correctly forecast the expected return securities and the expected return of portfolios. It basically comes down to our realistic probabilities, and those realistic probabilities incorporate the risk premium that essentially determines how these expected returns differ from the risk-free rate. What we really have is a joint stochastic process on everything; and we want everything to be the correct probability.

So in qualitative stylized facts, there are some properties of interest rate generators. I think Sarah has talked about a couple of them already. Interest rates don't go to zero and infinity; interest rates can spend up to several years within a narrow band of trading range and then the band changes. Long-term rates are correlated, but not perfectly. The volatility of long-term rates is less than short-term rates. Yield curves can have a variety of shapes, and the higher the level of rates, the higher the absolute level of interest rate volatility. So if rates are 14%, it's easier for them to move up to 14½% than if they're at three and move to 3½%.

So what do these facts tell us in terms of building our interest rate generator? By developing a list of all the qualitative stylized facts, we have a criterion from which to determine whether all our yield curve sequences are in the model. If we see that there are no yield curves that invert, we see it as a problem. We also want to keep those impossible yield curve sequences out of there because if you have impossible yield curve sequences, you might decide you need to hedge them. You may end up giving up expected return in order to have something that can't happen.

Obviously that would be pretty unfortunate. I just listed David Becker's facts. We looked at both the individual yields by themselves and also the relationship in terms of yield curve inversions. You can work out different points in the curve to get additional relationships.

What do these things do for us in terms of building our generators? A big problem with building our generator is we have only the historical path of yield curves. So how are we going to go from one path and figure out the probabilities of paths? What do we do? We say we want to build a model that's sort of robust, that has stable parameters, and that makes sense. Given that model, we can infer the parameters of it based on standard statistics. That gives us a way of developing probabilities for yield curve sequences, even though we just observed that one sequence. The quantitative stylized facts were used to make sure that we reproduce those with our calibration. We're also concentrating on the quantitative stylized-like facts that are most relevant to our cash-flow testing work. Basically that's going to be most sensitive to how likely high rates are and how likely yield curve inversions are. They're going to drive the results, particularly if we're looking at probability of ruin or just probability of low return. We use these to assess the model's forecasting probabilities of yield curve sequences, and we also can test whether we have all the possible yield curves in there. If, for example, we have no yield curve inversions, we'd know that some yield curves were just left out completely. We would also contest impossible yield curves. We have some statistics where events rarely happen. If those happen often in our scenarios, we know that we have a problem. Between these two sets of qualitative stylized facts, we can assess whether we have all the possible yield curves and have excluded the impossible ones. Quantitative stylized facts also help us assess the probabilities.

Now we're going to think about which interest rate generator we should have. We're looking at a variety of generators we could possibly have and we're deciding which one is a good one. And we start asking some questions. First, does the interest rate generator cover both the realistic probabilities and the risk-adjusted probabilities, or is it just one of them? Some products don't really address the fact that there are two different ones. There typically tends to be Wall Street based ones that just look at risk-adjusted probabilities because they're only interested in pricing. Sometimes they will encourage you to use those for realistic cash-flow testing, even though it's inappropriate. There's also the problem of using one common set of interest rate models. You always see those in risk-adjusted form and they're always the scenarios you use for pricing. They don’t address the issue of how to get realistic probabilities for those yield curves. So those models are particularly difficult to use when extracting the realistic probabilities. That makes them most difficult to use for portfolio analysis.

We've gone over whether it has both realistic and risk-adjusted probability, and whether it has one common set of yield curves and two different probabilities or whether it is two different sets of yield curves with different probabilities. Does the generator maintain a distinction between fixed parameters that are calibrated to history and random variables that are refit daily and simulated? Or does it really rely on refitting the parameters each day without simulating their change, in which case it can't possibly be doing a good job reproducing the universal yield curves. It's going to be leaving out a great deal. What I've said before is that the realistic probabilities are calibrated to a reasonable period of time.

Can the interest rate generator reproduce yield curves over a long time period with fixed parameters, or does it have to change its parameters all the time? If it's changing its parameters to fit each day, then when we generate yield curve sequences for cash-flow testing or pricing, we know it's going to be leaving scenarios out, and that gives us biased results. If we decide to use two different yield curve generators with different sets of curves for the two different probabilities, what sorts of problems is that going to give us? How different are the yield curve sequences in the two models? It just turns out to be too much work to have two different models that are different for the two different applications. It's much easier to have one model that's arbitrage free and that has both sets of probabilities and a common set of yield curves. Finally we must calibrate this and check Becker's stylized facts. The footnote here is that this combination of looking at Becker's quantitative stylized fact and yield curve inversion has been called Becker's razor. So the question is, does a generator pack Becker's razor?

We talked about the core interest rate generator. Now we have to talk about the whole interest rate generator system, because there's much more to the generator than just the core. There's also the question of its historical calibration. We need somebody to be doing that. We need to be having these parameters updated. If a vendor isn't doing it on the outside, can we do it internally? What are we going to have to do? We're going to have to develop many of the statistical codes, and we're going to have to maintain that and have somebody write it. It ends up being a great deal of work. Then there’s all this other information that outside vendors sort of develop over time. If we have to do that in house, it's going to be a real pain in the neck. Is the model sufficiently robust so that if we make a decision now based on analyses using the model, is that decision going to be something that's stable or is this re-parameterization process going to lead us to a different decision six months or a year later?

So what are the weaknesses of a one-factor model? This basically reviews them all. They can't do much of anything. They can't reproduce the qualitative stylized facts because it can't handle those different yield curve shapes and inversions properly, you can't calibrate them to Becker's stylized facts because there are too many to calibrate to. They can't match the yield curves over a long period, such as 25 years by just varying the one factor. They don't really identify expected return as the yield curve varies because in order to identify that expected return, you basically have to have the right probabilities. Since they don't have the right probabilities calibrated to the quantitative stylized facts, they can't correctly calculate expected return. They don't give you the risks of yield curve shapes, inversions, and sequences because they can't cover the full universe. The other problem they have is if they are mean reverting, rates will go up but they'll come back. So you can't get the effect where rates go up and just stay there or rates go down and stay there. If you're looking at your 3% guarantee, which is your floor, obviously you have not yet become concerned about the scenario where rates go down to 2% and stay there for five years. If your model is just a single mean averting process, your rate might go down to 2%, but then it's back up to 5% or 6%. It's basically saying that floors are free. It's something that might not turn out to be the case.

Basically the model doesn't represent the different possible events like yield curve inversions and so forth. If I'm pricing some complicated option that depends on the shape of the yield curve, then this one-factor model is not going to do a very good job. If it does a good job on one security, it probably doesn't do a good job on the next one.

So what are some two-factor models? The first two-factor model that people, primarily academics, used was the Brennan and Schwartz model. It was developed about 20 years ago. Two years ago, Hogan proved that not only was it a bad model, but that, technically, there was no solution to the model. They used the yield on a long-term bond on their second variable and assumed they could get away with that. It looked like they would, but then there turned out to be a problem.

Then we have the two-factor Cox, Ingersoll and Ross model. Say you have a short-term rate that is mean reverting to a target rate. The real rate and the inflation rate are mean reverting to a target rate. You can say that the interest rate is the sum of those two, or you can just say you just have two factors. You take the sum of those two, which is the nominal rate. The problem with this model is that the two factors can't be correlated, and of course they always are. So that's a major problem.

The next approach would be to take the short-term rate as our first factor, and make volatility the second factor. The problem with that is volatility doesn't have a big impact on the bond price. As far as the bond price is concerned, you have a one-factor model, which means you have a one-factor model for the yield curve. We just saw that one-factor models generating the yield curve don't work very well. If you're a Wall Street firm using the second factor of stochastic volatility, it may work well for you because you want to calibrate these curves to derivatives. You're only really concerned about buying and selling and laying off your risks. You're not in the same situation that an insurance company with liabilities on their books for 10 or 15 years is in. That insurance company is stuck with them.

There are these two-factor Heath, Jarrow and Morton (HJM) models that are extremely popular these days. The problem with these is that they're only in the risk-adjusted probability form. People just don't bother to work out what the realistic probabilities are. One of the reasons is because it's hard. You end up having to do all that yourself, and they're probably not even going to tell you that you have to do it. It ends up basically saying develop your own interest rate model. There's the double mean reverting process™ that I talked about at the 1995 annual meeting in Boston.

The other problem with HJM is that it's a very complicated way to say that I have parameters that I reset every day. They don't come out and say it that way, but that's essentially what happens. Technically HJM is really a methodology. When you take an HJM model that was originally put in an HJM form and then work it back to a state space model when that's possible, you often find that it's a silly model. This was done by HJM for the Ho-Lee model. It was also done for an HJM two-factor model. They basically projected what the state space representation of that model was about five years later and found that the model predicted that interest rates went to infinity with probability one. So that turned out to be a problem.

Similar things have come up with these other HJM models, but I think they finally found one that's reasonably stable, but it's still only the risk-adjusted probabilities. You don't get the realistic scenarios. If you want to use it for cash-flow testing, you're basically on your own. I think that for the insurance industry these are really of secondary importance.

Let's go over the double mean reverting process.™ If R is the short-term rate, which you can think of as a three-month bill rate, we take the logarithm to get this variable U, and then U follows this process. Now my λ_u is the market price risk for the U. There's actually another component that has gone to a new set of presentation, but this λ_u is the target rate, so you can think of U as turning towards a target rate but it doesn't get there all the time. It stays a distance λ_u away and then the theta goes to a grand target rate theta bar. This theta is essentially like a yield curve shape parameter because of the risk adjusted target rate. The parameter λ_u measures the average steepness of the yield curve and that's what is fit to reproduce Becker's yield curve and inversion frequency. What this model basically says is that U is attracted to theta, and theta is attracted to theta bar. It turns out kappa one is strong, like 0.5–0.7, and so U gets strongly attracted to theta minus λ_u. The theta determines the trading range and U is strongly attracted there. The parameter kappa two is low, so theta can move, but it doesn't have a strong tendency to go back to where it's going. If we're in the early 1980s, U can go up, theta goes up, U goes up with it, and basically the two of them can stay there for four or five years and then they slowly wander down because this kappa two eventually pulls theta down. Or we can have interest rates go down to 3%, 4%, or 5%, persist after three or four years and slowly we'll see the theta come back up. The yield curve will follow along behind it. You can plug in this S(T), and you can plug in a little residual curve, you basically fit U and theta to fit the yield curve as best you can, and then you plug in this low residual curve because some people absolutely have to fit the initial yield curve. This allows them to do that. To get the risk-adjusted form probabilities, we set λ_u and λ_θ to zero. We change kappa one and kappa two and that gives us the risk-adjusted process.

Let me just make a couple points about this model first before going on to the economic scenario generator. Basically, this one passes Becker's razor. It has been calibrated to quantitative facts; it reproduces the qualitative stylized facts. Rates aren't going to get a positive infinity or zero; they're going to be pulled back by the mean reversion. Because you have two factors, you can handle yield curve shapes. That has been tested in the past. With fixed parameters from 70 to 94, I could fit the different yield curves over the period just changing U and theta within 15 or 20 basis points. So that tells you it's representing the universe of yield curve shapes. We have long-term yields correlated with short-term yields, but it's not perfect. Because we're using logarithms, when interest rates are high, you get higher volatility in the *R*, so your absolute volatility goes up with the level of rates.

I think we've basically covered qualitative stylized facts. It's been calibrated to reproduce the Becker quantitative stylized facts, the frequency of yield curve inversions, the probabilities that different yields are in different ranges, such as between 10 and 11%, or 12 and 13%, and so forth. If you go back and check through all the other things and questions I asked for interest rates to satisfy, it just so happens to satisfy all of these requirements.

The next step we can look at is economic scenario generators. Given that we have a yield curve process that's reasonably stable and adequate, what is the next step? We want to add in some other variables to get an economic scenario generator. The first one we think of is inflation. For the last 30 years, economists have had the goal of building a generally clear equilibrium economic model, which based on inflation, real GNP, and unemployment, you could use to figure out what interest rates were. Basically that totally failed. Instead of building that sort of fundamental model, we had to build some statistical models. One approach to that is this little inflation regression, equation 0.3. It basically says *R* is the short-term rate, *S* is the exp (theta) and so it represents a target in sort of interest rate terms. *I* is the inflation rate. If short-term rates are high, that probably means inflation is high. If *S* minus *R* is a large number that means that *S* is the target rate, so that means that the yield curve is steep. If *S* minus *R* is low or negative, it's inverted. So if *R* is high and the yield curve is inverted, we're basically saying that *R* is high, so inflation is high, the yield curve is inverted, and that means the Federal Reserve did some tightening. It wouldn't be tightening unless there was high inflation, so that happens in the case of really high inflation. If short-term rates are low and the curve is really steep, that means the Federal Reserve is really being easy. That must also mean that inflation is low. So we can sort of infer inflation backwards from the yield curve, assuming that our Federal Reserve is operating in a rational fashion. There are stock indexes and dividend yield. If we want to do an equity-indexed annuity, one of our big problems is essentially that we have to look at both the interest rate and stock index together. A simple approach to that is if you buy a five-year zero and buy an option on the Standard and Poor's 500. The problem is, if interest rates go up, bond prices go down, and the stock market goes down. Your zero-coupon bond goes down at the same time that people are surrendering, so you end up having a loss, if that was your simple strategy. So to address those sort of questions, you must have a joint model of the interest rate process and a stock process. You also need the dividend yield in there because the way these equity-indexed annuities are structured you only get the price return. You have to take the dividend out, and obviously there are other applications beyond that particular one. It's a fairly hot one these days.

Mr. Steven P. Miller: I can produce interest rate scenarios that do most of those things and in some cases all of those things. However I can't tell anybody what the probability is that next year's interest rates are going to be 6% or 9%. In other words I don't know what a realistic probability is. I do agree with you that the probability that I arrive at an arbitrage-free scenario isn't realistic but knowing that something isn't realistic is different than knowing what is realistic.

Mr. Tenney: Basically it's a process of figuring out market price risk which lets you go from one probability set to the other set. If you set up your risk-adjusted process and you work with that and now you want to get some realistic probabilities, essentially you're going to have to introduce some risk-premium coefficients so each of your variables can change randomly. You're going to have to estimate those risk-premium coefficients to reproduce like Becker's razor does. You're saying the expected return on any security equals the risk-free rate plus essentially the elasticity or the duration with respect to each of the variables times the volatility of your variable times your risk-premium coefficient. You add a term like that for each random variable you have. Then you find those parameters that give you realistic probabilities.

Mr. Davlin: A question from the audience. Was the distinction between state variables and parameters in these models clear in everyone's mind? Would it help if Mark clarified that a little bit?

Mr. Tenney: Well, basically, state variables are what are changed randomly from day to day and parameters are what stay fixed. So if you generate your scenarios, the parameters are the numbers that don't change, and the variables are the ones that do.

Mr. Davlin: So one major criticism of the one-factor models is that they're treating the parameters like variables.

Mr. Tenney: Right. When it's convenient they treat it one way, and when it's not convenient they treat it the other way.