RECORD, Volume 31, No. 1*

2005 New Orleans Life Spring Meeting
New Orleans, Louisiana
May 23–24, 2005

Session 30PD
Setting Long-Term Investment Assumptions for Actuarial Models

Track: Investment, Pension

Moderator: PHILLIP SCHECHTER

Panelists: KEVIN C. AHLGRIM
PHILLIP SCHECHTER

Summary: Actuaries are increasingly involved in the investment assumption-setting process for multiple purposes-pricing, hedging and long-term strategy setting. As anyone who has performed this exercise knows, it is a combination of "art" and "science." This session discusses the process of setting long-term investment assumptions from a practical perspective. Topics include: Has the equity return premium changed? Does it matter? When to use risk-neutral scenarios versus "real-world scenarios"; methodology and considerations for generating risk-neutral scenarios; and special considerations for setting pricing scenarios.

MR. PHILLIP SCHECHTER: I currently work for AXA Equitable in valuations and forecast. I started with Mutual of New York in 1987 and stayed there until AXA purchased us. I've been involved with modeling net investment income and producing asset models for cash-flow testing since 1990.

Also presenting today is Dr. Kevin Ahlgrim. He's an assistant professor in the Department of Finance, Insurance and Law at Illinois State University. He has his B.S. in actuarial science and an M.S. and Ph.D. in finance from the University of Illinois at Urbana, Champaign. He teaches on the subjects of investments and

NOTE: The chart(s) referred to in the text can be downloaded at: http://handouts.soa.org/conted/cearchive/NewOrleans-May05/030_bk_new.pdf.

financial risks in management, and he's been involved in producing Web-based courses on financial risk management for insurance enterprises, which is good for continuing education credit, if you're interested. His interests lie in application of fixed income techniques to insurance. He's worked as a health actuary for CIGNA and AON.

When I think about investment assumptions, I think about the work we've been doing for almost decades now on cash-flow testing in which we've been looking at deterministic scenarios. We've been looking at spreads of treasuries, default rates, prepayment rates, asset expenses, real estate returns as a general asset class and equity returns, maybe related to interest rate scenarios, and maybe just as a separate test on the side. There are also stochastic scenarios, either arbitrage-free or real world in our models. What can we expect from future returns? How should asset classes be related? This is certainly the direction we're expanding into with financial economics, but we still have to make assumptions about a number of items. I'll begin by talking about some of these.

Traditionally, we've set default assumptions by looking at long-term averages of each asset class. For publicly traded asset classes, there's been a fair number of good outside studies—Moody's for instance where they rate issues and look at histories. We try to come up with one number for asset class that expresses default cost, which is really a combination of default probability, write down on asset loss during a nonproductive holding period and any residual value at the end. We reduce that to one number usually because it's much easier to deal with and has a somewhat more stable pattern.

However, reducing our default cost to one number has made it comparable to another asset assumption we're making, which is spreads. If we're taking a look at a AA bond, we assume going forward a certain spread relative to treasuries and we assume a default cost, and that leads to the concept of net default cost and net spread. Being efficient frontier types, we believe that the more risk we take on, the higher our expected return. So as you go to riskier asset categories, you're expected to pick up additional spread quicker than you pick up additional default costs. This could lead to a situation in your models where you're just modeling by this one number, this net spread, where the riskier assets you go to the better your results look.

New York recognized that issue in regard to cash-flow testing. They said, "Okay, you have to run at least a sensitivity where you're not picking up any more than 100 basis points on spread, which is a good regulatory type of an approach. That way you can compare companies and see how much they are picking up in spread."

But it really ignores the root economic problem of this analysis, which is that we're using a deterministic model to model a variable and apparently random cost. It's not the expected cost of defaults that's killing us; it's a bunch of defaults at once. So there are two ways we can improve our default models, and those two ways could be using stochastic defaults and/or using numbers that change by scenario.

I'm going to assume we're modeling asset by asset. And with stochastic defaults, let's say you expect your default cost for this category to be 25 basis points a year. Instead of decreasing your yield or decrementing your asset by 25 basis points, toss a loaded coin. If it comes up heads, you're safe. If it comes up tails, the asset is in default with no residual value. Running this type of an analysis a few times will give you some measure of some variability of results. It will capture some of the risk of default, and it will also capture the fact, although not very frequently, that the rate of default can be quite high, as we've seen a few too many times this century already.

A second related approach is to have defaults that are dependent on other elements in your scenario. What you want to relate your default cost to would depend partially on what's in your scenario. But basically, you want to capture items that influence the rate of default. So, for instance, if your gross domestic product (GDP) is going nicely, you'd use low default cost. If your GDP is flat or decreasing, the default cost would increase. Unemployment could be a proxy for economic health. If stiff equity market returns are going very poorly, you'd expect to have more defaults. If you have a fairly old-fashioned model and if you're not a separate account company ,your scenarios are really all framed in terms of interest rates. Given the New York Seven, you could assume that big jumps in interest rates are caused by or will cause uncertainty and increased defaults.

If credit spreads are increasing, you might say that the probability or expectation of defaults is increasing. I think that's a fair statement, but there have been times, especially in the past three years where, even as the outlook for default seems to be increasing, credit spreads were decreasing. This was probably because the stock market wasn't doing well; interest rates were very low. All of us were desperate to pick up spread any way we could, and because of supply and demand we were willing to pay for spread. We were willing to pay for yield, and companies with extra risk were able to issue cheaply just because we were desperate for any yield we could find. So it is probably true with equity market returns or interest rate jumps.

There are more things going on than just defaults, but it would seem that defaults should somehow be related to these items. I think some combination of stochastic defaults and scenario definite defaults could give you models that are possibly more useful and definitely more interesting than just hitting every asset with a flat yield haircut.

Another assumption that we have to look at is prepayments. When we started doing cash-flow testing in the late '80s, prepayments took a lot of our time and energy because we were in a time of decreasing interest rates and we had portfolios of bonds with either no call penalty or fixed call penalties. In the 10th year, you prepay for a 3 percent penalty, in the 11th for a 2 percent penalty and the like. In a decreasing interest rate environment it was advantageous for them to prepay and they did. So we were hit with collecting large pools of cash in a decreasing interest rate environment, and reacted by including this in our modeling. We developed models that looked at the market value of the bond if it were not prepayable, compared to the call pricing, assuming that if the market price of those cash flows is higher than the call price by a certain margin the bond will be prepaid.

Now we're in a slightly different environment, partially because of what we demand from borrowers, most of our publicly traded bonds and even our private placements and our mortgage loans have some form of make-whole provision. A make-whole provision means that the prepayment amount is pretty much the market value of the promised cash flow. So it's really no economic loss because you're paying the value of the higher cash flows that you originally promised. This, from an economic viewpoint, in a frictionless world can make prepays pretty irrelevant.

However, I want to bring up a couple of points on prepayments. Can we have more prepayments now that most assets are yield protected? The answer is yes, but there are issues. The first is a data issue. If you have a commercial mortgage portfolio and you make a quick phone call to your asset manager and say, do I have yield protection on this, the answer is probably going to be yes. If you actually look at the terms of the loan, especially if it's a more complex structure like a construction loan or some other loan that has resets built in, you might find that there are periods during the term of the loan where you can prepay either whole or in part without any penalty or with minor penalties. These types of assets aren't captured well in accounting systems. It's very hard to get these multipart loans in a standardized format. For a lot of companies the mortgage loan portfolio is really dominated by a few large loans. It's really worth spending the time to talk to your asset managers who might be able to model your top 10 or 15 loans that could very easily be 50 percent of your portfolio. It's worth understanding those.

The other issue we have is with these prepayments coming in at a time when it's very difficult to predict because interest rates are not driving them. It's very often corporate restructure and the like. Your investment department may be looking to you, as they frequently do, to determine cash flows in both liabilities and assets for the coming year. You could give them one number and suddenly be faced with 50 or 100 million in prepayments. You can tell them that the prepayments weren't in your model, but they're stuck with the cash they have to invest. So you should warn them upfront if you're ignoring prepayments.

If your model has assumed very precise portfolio management strategies where you're holding certain types of assets and very tight ratios, it's fair to assume that if a large prepayment comes in, your investment department is not going to be able to turn it around very quickly and invest the remainder of that asset, even if they recognize where the cash came from. There is also the impact on expenses, which translates into higher transaction expenses.

One asset class is an exception to this type of logic, and that's asset-backed securities (ABS), and particularly mortgage-backed securities. If you ever try to trace them through manually, they're crazy. You have these collateralized mortgage obligations (CMOs) where you'll have a hundred different trenches backed by dozens and dozens of mortgage pools in various combinations. Now we can have other collateralized assets as collateral. It's very tough to really look through to the collateral and make your own assumptions about what's happening. Consequently, we're using more external models. Intex and BondEdge might provide data on the yield structure. Davidson might provide software to use to look at each individual pool and assess its risk of prepayment.

These models are great. But I would say, having been raked over the coals by auditors before, the fact that we don't effectively use these models doesn't get us away from our basic responsibilities to ensure that the outputs are reasonable. The ways we can do this are by looking through the underlying formulas on prepayment logic and particularly on individual mortgage pools on pass-throughs. Let's take a look at what we've been forecasting for the past couple of years and compare that with what's actually happened, because these things are cutting down each period and there's real data there.

The backdrop for everything is the sensitivity test. One cautionary tale is what happened when we started getting into CMOs back to the '80s. We had a good 20 years of experience studies on how homeowners prepay mortgages, and we thought we understood that pretty well and, in fact, we did, but then this whole asset class took off. There were mortgage pass-throughs. There were CMOs backed by Fannie Maes and Ginny Maes and all these secure assets. The secondary market really took off, and it paid for a bank to issue a mortgage, sell it immediately as a secondary market and make a buck.

What happened is more banks became interested in doing that and you started having ads on TV and in newspapers and radio, refinance your home. The process for refinancing got easier and easier. So all our experience studies about how homeowners behave went out the window and we, as an industry, and everybody invested in these things, had bought significant numbers of the securities. Because the yields were pretty good, we were paying a premium over par for these assets. These assets were running off the books at very fast rates and we were getting, in many cases, negative returns. We learned from that experience and it became axiomatic that you never bought a CMO trench at a premium. Now things have stabilized enough that we don't have to buy it at a premium because we think we understand the behavior in a new world.

One concern for testing, is if interest rates go up, there's an extension risk, which I don't think we have paid much attention to in the past 10 or 15 years. The commercial models seem to assume there's a certain floor below which the level of prepays will not drop because of people moving away. But it just seems to me that people have this option. If you have a 4.5 percent mortgage and interest rates go up to 8 or 9 percent, even if you move, you're going to try to hang on to this asset in some way. You're going to try not to prepay if you have money at below market rates.

If there are hundreds of thousands of people looking to not pay off their mortgages, there are going to be legal structures or financial structures coming into play to allow people to not prepay. Something where you don't sell your house, you rent it to somebody and they get some secondary equity that way. I don't know what the structure would be, but if we're handing out options to people, you have to assume that people are going to take advantage of them because that makes sense.

Now I will talk about the new exciting financial-modeling-type assumptions. Has the equity risk premium changed and does it matter? I'll answer the second part of that question first. Yes, it matters. The equity risk premium is the expectation of return on the part of the investor minus the risk free rate. In other words, how much more do you expect to make in the stock market than investing in treasuries? This is basically a way of asking how much we are going to earn in the market.

There are a couple of sides to this argument, and I'm going to mention Robert Arnott, because if you search the Web on this question, you'll find a lot of articles, and most of the yes answers end up being attributed to him in one way or another. He has this decomposition of return on stocks where what you expect to earn on the stock is your dividend yield plus your growth in value. Your growth in value can be broken down into growth in earnings plus the growth in the price earnings multiple change. If earnings double, you've doubled your price.

Historical dividend yield has probably been, depending upon how you look at it, around 4.5 percent. Now it's less than 1.5 percent. Growth, historically, has been around 5 percent. I don't know the future. Early last century and well into the middle of last century pricing earnings multiples hovered around a factor of 10. Now they're up to 30, 35 or 40. They were up over 40. I think they've come back down. It seems logical to say we won't see an infinite increase. At some point, people aren't going to pay 100 times the value of earnings for a stock, on average. So if you look at it this way, you'd have to say, yes, we can't expect the type of returns in the market that we've seen in the past.

On the flip side, we believe in the efficient market hypothesis for the most part. Anything that we know, as individuals, about the price of a stock has already been incorporated into the price. So the fact that I've rattled off some numbers and statistics is irrelevant because the market already knows this. So if the P/E ratios are high, that could be because we expect future growth. If yields are low, the companies are doing better things with the money. They're hoarding it for growth. They're deploying it. So efficient market people, I think, have to be reluctant to say we know what direction the market is going to go relative to what it's been doing.

This is all complicated by the fact that you can't measure the equity risk premium directly. The typical person, the guy in the next cubicle trading all day, isn't thinking about what rate to expect to earn on this investment relative to the risk-free rate. Probably if you ask your investment department the same question, you wouldn't get much better an answer. They're there to find cheaply valued stocks and to buy them. So people aren't directly thinking of this number when they invest. The current estimates range from 0 to 5 percent. You can have a wide range of current estimates and a number of ways to try to get an estimate of value.

I think the efficient market people win. When you look at our actuarial literature that deals with generating scenarios, for example, one of the first things done is you develop your model, and then you set the parameters to match history. This is true of the risk-based capital (RBC) working group. This is true of the Mary Hardy article on regime switching models. I think it's true of the set of parameters we'll see later today. The first thing you shoot for is to reproduce history. Of course, that has the ancillary benefit that it lends to model credibility. That can produce something that looks like history. So you can use them as provided or you can reduce it somewhat to reflect the valuation ratios are currently high or whatever reason you think the market might go down.

Ultimately, it comes down to how you are using the results. If you're using the RBC committee, where you really don't care about the average, you're throwing away 65 or 90 percent of the scenarios and you're just looking at the tails. So there the average is probably not what's really important. It's really what the tails look like. However, if you're pricing or doing corporate forecasting, you'd rather look at the averages. You don't want to go to management and say, I came up with a new product and I think there's only a 3 percent chance that it's going to wipe out the company. I don't think that would exactly get approval.

Now we're going to talk a lot more about stochastic scenarios. Why the interest in stochastic models now? We've been talking about it for a decade as a profession, and now we're actually using them. What's driving us? I think there are two major points. One is that we're selling products where the path of market returns matter as much as where the market is in five, 10, 15 or 20 years. We have guarantees that could be reset every year. We really care about the path, and it's hard to come up with in a deterministic way what the points and path are that you have to look at.

Next with stochastic scenarios, you have some regulatory and professional issues. SOP-03-01 doesn't explicitly mention stochastic test, but I think most of us have interpreted it to mean you have to do stochastic testing. And RBC C-3 certainly does mention it. But these aren't what are causing the need. I think they're responding to the need and trying to get a coordinated professional response. The real need is the sensitivity of the variable products.

Another major point is technology. We have really fast computers. We can do the modeling now whereas when we started talking it was more theoretical.

My next point is: Are we going to use risk neutral scenarios versus real world? Risk neutral scenarios would mean you're not compensated, at least on an expected value basis, for taking on additional risk. So you'd expect, by definition, the equity risk premium I was talking about before would be zero. Any additional spread you've got from investing in the risk asset categories, you directly offset by default cost. It doesn't sound very realistic, but it's consistent with the underlying option pricing models and it's also consistent with the fact that a dollar today is a dollar today no matter what future yields are. So if you're just reproducing what you're holding today and you want to make it consistent with the options market, and have market price optionality, I think risk neutral has a strong advantage.

Real world scenarios is a way to say we expect more from stocks than from risk-free bonds, and it's really a crucial key. For instance, we act as if we believe in the efficient frontier. The efficient frontier would not be very exciting if the expected return for everything was the same. If done correctly, it will produce the price of the currently existing options. If you're setting economic capital based on this approach, your current economic capital is probably higher if you invest in a riskier class of assets today. So you're realizing investing a dollar today is producing two different values today.

It does lead to some conflict because you're saying, I'm going to earn more on this, and therefore the economic value of my company is higher right now. But overall, for producing future cash flows, real world is probably a better way to go. This is basically a set of discount rates, which lets you use the real world kind of scenarios to project, but discount in such a way that you produce current market-consistent prices.

Next, I'm going to breeze through scenario generator parameters. This really is just a matter of terminology between geometric vs. arithmetic mean. You have to be careful with the terminology. People ask what should the mean be or what was the mean of this set of returns? It's pretty basic stuff. The arithmetic mean is going to be different than your geometric mean. And the bigger the variability, the higher this difference is going to be. So, for example, if you lose 1 percent one year, gain 1 percent the next year, your arithmetic average is you've had 0 percent returns. On a geometric basis, say you're down to 99, and then you're up to 99.9. You take the square root of that and you have a slight negative return.

A more dramatic example is if you lose 50 percent and then gain 50 percent back. Your arithmetic average is still zero, but your geometric average is very different. You've gone down to 50 and up to 75. The square of 75 is 0.86, so you have lost 13 to 14 percent a year. The bigger your variance, the bigger this difference will be. It's just a matter of terminology and being careful. You're talking about making sure that when you're communicating with somebody you're communicating the same thing. If you're assessing risk, you care about the tails. If you're trying to say what's going to happen, the mean is pretty important.

Here I just ran a fictional set of DVIV guarantees through a bunch of scenarios all with the same averages. One was a regime switching lognormal approach. The second was just one scenario at the same average rates. The third was using lognormal scenarios with 15 percent annual variability with the same average rates. The upshot is more or less the same pattern. Years ago when we started talking about stochastic scenarios, my first thought was, you have a way of making probalistic statements and that makes your individual assumptions less important. As I started working with them, I realized, no, it just fits the load of the assumptions to another place. But changing your assumptions as to how you come up with your scenarios can still be as important as how you pick an individual scenario.

Before I had mentioned that your default should be consistent with other items in the scenarios, and that's probably true of all items in the scenarios. There should be some consistency, just like Reg 126 required that your prepayments were consistent with your interest rate assumptions. But now, going into scenarios, it's becoming a bit more complicated to say what should be related and what sort of economic scenarios are consistent. This was a concern both for the Casualty Society and the Society of Actuaries and they're sponsoring a study on this issue. We're lucky to have Dr. Ahlgrim here with us to discuss that study.

DR. KEVIN AHLGRIM: I'm going to introduce you to a research project that was initiated jointly by the Casualty Actuarial Society and the Society of Actuaries. It has a quite long title, "Modeling of Economic Series Coordinated with Interest Rate Scenarios." My colleagues that worked on the project with me are Steve D'Arcy, who is now the current president of the Casualty Actuarial Society, and Rick Gorvett, both of them at the University of Illinois.

I could have put in a second disclaimer that says, of course, this is not necessarily the gospel according to the Society of Actuaries. It just represents our interpretation, but they provided some very helpful guidance and feedback on this subject.

Essentially, what I'd like to do is talk about that model; provide some of the motivation for that model; give you a short overview of the model; discuss some of the series that are modeled; and describe some of the underlying mechanics and the map behind the processes. I'll talk briefly about some of the applications of that model and also compare the results of that CAS/SOA model, which I'm going to call the financial scenario generator, to the infamous C-3, Phase II 10,000 scenarios that are on the Academy's Web site.

In 2001, this research was initiated jointly by the two actuarial societies here in the United States as basically a tool that actuaries could use for many of the applications that they get involved in and many of the analysis that they look at. Many of the different kinds of analysis that actuaries use have these base underlying assumptions, such as what happens to stock returns, what happens to interest rates and things like that. So they were interested in, perhaps, coming out with a project that would help generate some of the financial scenarios.

Let’s talk about some of the different areas. Of course, on the casualty side, you might consider dynamic financial analysis, which is just cash-flow testing, or solvency testing. It could be used for pricing or reserve setting. Many different applications use similar kinds of assumptions. So the goal was to provide a tool for actuaries to set some of those assumptions and to look at some different scenarios with those kinds of economic and financial barriers using realistic interdependency among those variables.

It's important to point out that this, we believe, is the beginning. It's the foundation of what we hope is an ongoing dialogue with the industry. We don't expect it to be the last word in economic scenario generators, but we wanted to start out with a basic foundation to move forward that generates a discussion in the industry among practitioners and get feedback to improve upon that over time.

The scope of the project was basically three-pronged. We started out performing a literature review, and looking at some of the major papers that have been proposed for different kinds of financial series from finance, economics and actuarial science. The second was the creation of a model itself, which is a spreadsheet-based model. You can download that model and begin to play with that model. I'll talk a little more about that later, as well as provide documentation of that model. That documentation is available on the Web site, at http://casact.org/research/econ/. The third was the report. The report structure currently has the first section basically being an introduction to and overview of the problems and motivation for the study. The second section contains excerpts from the RFP and the third the selected proposal. The fourth is a little talk about the literature review—summaries of some of the major papers that helped frame the work that we did.

Section five actually talks about the approach that we used to model different kinds of series. We'll look at some of those shortly and talk about the data we used to choose the parameters. We wanted to use publicly available data so that our work could be replicated.

Section six talks about different issues. These are more of the enthusiastic discussions we had with members from both societies about the way they would like things to be done. We had a lot of discussion about what interest rate model is the best or what's the best way to approach this subject. These are basically our responses to some of those issues that were generated during that feedback.

Then we run the model very similarly to what the C-3, Phase II scenarios do, and the discussion of those results is in section seven of our report.

Then there are some appendices. Appendix A actually describes how to use the model. It describes the spreadsheet itself, and gives the overview of the spreadsheet. It does use an at-risk simulation add-on. If you haven't used at-risk, it's a very powerful simulation tool. If you get involved in any kind of stochastic scenario generation, the output alone in at-risk makes it worth the purchase. It really allows you to track virtually any kind of variable that you're modeling in a spreadsheet. It can capture that and put it into a spreadsheet or a lot of different graphical forms and provide a lot of different output.

So the user's guide sort of walks the user through using the model, but also talks about some of the parameters and how important those parameters are as you change them a little bit. It addresses what kinds of things are affecting the spreadsheet; this is how much interest rates are going to change, and so on. The next appendix contains presentations of this research, including this presentation that you see here today.

Appendix C is the C-3, Phase II. We have downloaded and dumped 5,000 different scenarios. And later on, I'll actually compare our model to the Academy Web site.

Appendix D is the model itself. You want to download and play with it. Absolutely do it. It's open and available for everybody. It's for your use and for your analysis.

Some of the work that's been done before in the area that you may have heard of is the Wilkie model. It's probably used more internationally than in the U.S. It really starts with inflation as sort of a core variable and then begins to incorporate other variables in the model. Hibbert, Mowbray and Turnbull provided some more modern financial techniques to a very similar kind of model. For our project, we combined those two approaches and applied it to the United States.

There are a lot of different models that are available. How do you decide which one is best for you? Of course, the main thing is to try to understand what your application is and how precise you need to be. Really, it comes down to a tradeoff between precision and accuracy versus the complexity of the model.

For pricing applications, especially if they're short-term pricing applications, many would argue that the accuracy demand is much higher. But if you're just doing a strategic planning exercise, maybe accuracy isn't necessarily that important. Some stochastic generation may approximate what can happen in the future, but the combination of so many other variables makes it very difficult to pinpoint the 99.99 percent of the distribution. The point is, if you're doing things where a whole bunch of other assumptions are going to come into play and you don't have a lot of data to base your decision on, maybe a more complex model doesn't add much value. There's a tradeoff.

I make this warning because whenever we came across this complexity versus accuracy, more times than not we chose the simpler model. The reason is we want to be able to communicate this with the whole industry and get more and more feedback and to see how well it actually performs and whether or not it's realistic.

There are better models, I will admit, in terms of interest rate models, but how much do they actually buy you? We don't think as much based on the additional complexity that's necessary to incorporate in the model. It would be too difficult to explain to the entire industry, in my opinion.

Ahlgrim Slide 8 shows some of the economic series that are in the model, and the relationship between these different series. Rather than choose any specific variable as kind of a centerpiece of the entire model, we really are combining inflation and real interest rates.

Ahlgrim Slide 9 shows how we modeled inflation. The parameter kappa is the mean reversion speed. If you want the variable to go back to that level very quickly, Kappa is high. If you want it to revert more slowly, kappa is low.

In addition to the expected change over the next instant, there's volatility. Of course, inflation is uncertain. That's what the second term represents. It's a fancy way to say take a draw from a zero-one normal random variable, but you want to scale it. How much volatility do you want? You can scale that by the volatility parameter sigma. If you want a very volatile process, make sigma really big and inflation will bump around all over.

In Ahlgrim Slide 11, it is safe to say that real interest rates are just slightly more complicated than the inflation process. The problem with using a one-factor interest rate model, similar to what we just saw with inflation, is that if one stochastic factor affects all interest rates. The result is that yields of all maturities, whether one-, 10- or 30–year, are perfectly correlated. As one goes up, all other interest rates always go up. Not necessarily by the same amount, but in a predictable way. So we wanted to allow some flexibility. Here's one part where we decided we need a more complex model. We looked at the two factor Vasicek term structure model. It's very similar to what we saw with inflation.

Here the short rate is mean reverted. By the term short rate we mean a very short interest rate. It reverts to some mean reversion level *l,* so you can think of this as a very long-term rate. But that mean reversion level itself is also stochastic. So the short rate is reverting to a long rate, which itself is stochastic. But the long rate reverts to some long-term mean reversion level as well. So now we have two factors that are impacting interest rates, basically, short-term interest rates and long-term interest rates. And these can be correlated, which they are in practice. When short rates go up, long rates usually go up as well, but they don't have to. It's an imperfect correlation between these draws of the random variable.

Nominal interest rates then just combine those two effects; inflation and the real interest rate process, as can be seen in Ahlgrim Slide 12.

Ahlgrim Slide 13 basically shows how the model compares to history. The model actually allows negative interest rates, but there's also an option in the model if you don't like that feature. You can set it on a lower bound. You can also set a lower bound for inflation. But in this particular run, we just looked at the full distribution. We looked at this and thought it's not so bad compared to history. There are also pictures of long-term rates in the report.

Now I will discuss equity returns. I've heard several times in different presentations that many people use the Hardy's 2001 North American Actuarial Journal paper for regime switching. The idea is that, empirically, if we look at equity returns, there seem to be fat tails. A good specific case of that is October 1987. We were many standard deviations out on the distribution, based on what happened in October of 1987, it seems almost impossible. So what we do is we introduce this regime switching model where returns come from two regimes: a high-volatility regime, which may have been where we were at in 1987, and a low-volatility regime. You can switch between the regimes at any point in time.

We model the excess returns. It's a little bit different than Hardy did. We look at the excess return over the risk free rate. Remember, the risk free rate is inflation plus the real interest rate. That gives us the nominal risk free rate plus some excess return, basically a risk premium for stock. We do this for small stocks and we also do this for large stocks. This is really not a lot different from C-3, Phase II. They use a regime switching model. They define their regimes a little bit differently, but in the end the results actually turn out to be very similar.

Ahlgrim Slide 15 shows one case where we actually did use some private information. We went back to 1871. This shows the actual history versus the model, which, of course, is a little bit smoother. But you can see we get a little fatter tails than the normal distribution by itself would assume.

Ahlgrim Slide 16 shows a similar distribution for small stocks, although we didn't have as long of a history for that.

Other series are variations on a similar theme. In Ahlgrim Slide 17 we use something very similar to the Ornstein-Uhlenbeck process for dividend yields and real estate. By the way, we were not able to find a really good historical series for real estate that's publicly available. So if your company has a large portfolio in real estate, you may want to use your own, perhaps, proprietary model. The model allows you to plug in pieces that you'd like to change like unemployment.

Once you've narrowed down the kinds of models and you've schooled yourself on different kind of interest rate and stock return models and figured out which ones you think are best, you're still not done. You still have to choose the parameters for those models. How can you choose the parameters? Well, there are basically three different approaches that you can use.

The first is just to use judgment. Obviously, that is not the most scientific way, but as you use the model and do some projections, what does the output look like?

Does it seem reasonable? That could be one way that you might actually choose parameters. Perhaps, a better way is to look at statistical estimation. Let's use the one sample path we have from history for interest rates. For equity returns we try to derive the parameters that best fit that one sample path. The difficulty here is that you must jointly estimate all the parameters. Now, what does that mean?

Well, consider a mean reverting interest rate model. If we have a mean reverting interest rate model, some of the volatility that we see in interest rates is because interest rates are reverting to some long-term average. That's not volatility. So we want to factor out the portion that is attributed to mean reversion to estimate our sigma parameter, our volatility parameter. We have to factor out all that noise first. That's what we mean by joint estimation. You have to do it all at the same time. You can't simply look at what's happened with the change in interest rates over time and call that your volatility or you're going to overestimate your volatility. Joint estimation, in some cases, depending on the model, can be very difficult. It's not as straightforward as just running a regression. Some models you can. Ornstein-Uhlenbeck allows you to do that. It's more straightforward than some of the other models.

So we have judgment and we have historical estimation. The third approach would be calibration, and that is to choose the parameters to best fit market data. This is what a lot of investment bankers tend to do. Suppose we use this model to price some set of securities. Ideally we'd match perfectly with what the market prices of those assets are, if it's a pretty liquid market.

Of course, it's difficult to match prices for insurance liabilities, but for each kind of similar options or similar securities we want to try to match it as best we can. You could use various measures of fit, such as least square, or more sophisticated approaches. We primarily used history, part of which is because we wanted people to be able to use the model to replicate our results. But, you can, of course, change the parameters as you see fit. If you'd rather use calibration, go ahead and use different parameters.

The model itself is an Excel spreadsheet. It uses the @Risk add-in. Many people have asked us why we don't come up with a stand-alone product rather than have an @Risk add-in. My answer is that if you're serious about stochastic generation and stochastic scenarios, you should probably have @Risk anyway. Otherwise, you can use the data dump that we have of 5,000 different scenarios, or you can use the Academy's 10,000 scenarios.

We go through 50 years of projections and we're going to project all the different variables that we just talked about. So we look at 50 years worth of interest rates. In fact, you get the whole yield curve. So unlike the Academy's model where you're only getting a few points, you get the entire yield curve: small stocks, large stocks, inflation, unemployment, real estate, things like that.

Stochastic simulation is obviously useful for demonstrating solvency; defining economic capital; proposing hedging solutions; financial risk management; reinsurance; and anything like that.

In Ahlgrim Slide 23 I wanted to actually show how our model compares to the Academy's. Their focus was on interest rate risk and equity risk, of course, for variable annuities. I'm not trying to say that one model is better than the other. I simply want to make you aware of the differences in the models so that you can use your own judgment as to whether you think for your application one is more appropriate than the other.

It's really hard to actually see what's going on with the 3-D graph. The difference between the financial scenario generator and the C-3 Phase II scenarios is difficult to see.

Ahlgrim Slide 24 takes a different slice of that distribution so we can get some sense for what's going on. This shows the three month nominal interest rate. This is basically output from the @Risk. One line is the mean; another shows +/- 25 percent, and the third shows the first and the 99th percentile. So we're just giving you different percentiles on the distribution. What you can immediately see is that interest rates are awfully tight in those scenarios. You can see that at first we're actually showing it by month and then after that we show it by year. So the kink is because we're going from monthly to yearly.

But you can clearly see that the volatility assumed in these two models is very different. Again, you can always change the parameters of either model. In fact, you're given the tools to put in your own parameters as you see fit. But based on the default parameters that are in each model, it looks like the financial scenario generator has a wider dispersion of different interest rates. You can see that more dramatically by looking at the histogram in Ahlgrim Slide 25. So you can see a very spiked distribution for the AAA model. We put a minimum of zero for the interest rate. That's why you see a little bit of a spike here for the CAS/SOA model.

Perhaps this is a little unfair because we're looking and then we compare it to actual data. Well, the Academy's model is really comparing it starting today versus history. Somebody suggested in a prior discussion that maybe we should look at the change in interest rates but, nonetheless, simply by looking at this volatility you can see there's a major difference in the volatility assumed between the two models. A similar thing happens for 10-year interest rates, as seen in Ahlgrim Slide 26. Again, there is a much tighter distribution for C-3, Phase II than under these financial scenario generators. I'm not trying to point out shortfalls of either model, but rather just making you aware of an issue related to the models. Again, looking at the histogram in Ahlgrim Slide 27, the histogram is much more spiked.

When you look at stock data, they use a very similar approach. They use the regime switching approach. The parameters look different, but, in fact, if you go through the math, it turns out that the average return and average volatility are roughly the same between the two models. In Ahlgrim Slide 28 you can see very little difference among all of them. In my opinion, they both do very well, at least when compared to history. This goes back to the equity premium issue though. Both of these assume that they're very similar to history. Maybe you would argue that America or the United States is generally richer and people are less risk-averse than they were before. They're more willing to buy stocks and so maybe the equity risk premium has come down. That's a valid argument. The same thing is true for small stocks, as seen in Ahlgrim Slide 29. You see pretty similar results when you compare the two models. In fact, it looks like the Academy's model shows a little bit higher risk premium than the financial scenario generator.

We look at this as sort of a beginning of an ongoing dialogue among CAS members, SOA members and other practitioners to try to get some sense for improving this model. This is supposed to be a public access model that people can look at, provide feedback on and use for their own specific purposes.

MR. LARRY GORSKI: I'm chair of the Academy group that you referenced a couple of times. I have a question and a comment.

First, I want to address your references to the Academy prepackage scenarios. You are aware that the prepackage scenarios were changed about five or six weeks ago. They are no longer driven by a regime-switching lognormal model, but a stochastic log volatility model. So the comparisons you presented were, I think, based on the prior prepackage scenarios.

Secondly, the interest rate generator used for the new prepackage interest rate path was based on the C-3, Phase I model. The old stochastic model was dropped.

I have downloaded all the material you refer to and have the @Risk package. It really is excellent. Anyone who's interested in this should get the material and get a trial version of @Risk just to see how it works.

MR. AHLGRIM: That's a good point. There is a 30-day sample period that you can use @Risk to see whether or not you like it. It's fairly easy. More advanced features, of course, would take a little bit more time to get into, but you can at least get the sampling product.

MR. GORSKI: It's 30 days or 100 simulations.

But in terms of maybe next steps, one of the things that have come out of this whole Academy process is a need to deal with the correlation between funds in a more sophisticated manner. And there's been some talk about replacing the linear correlation coefficient approach with the use of populism, and I'm wondering if you have any suggestions on that.

MR. AHLGRIM: I plead some ignorance. I actually haven't really used populism a lot.

FROM THE FLOOR: The basic benefit is it allows you to study the marginal separately from the dependency structure and the different models that exist allow you to have stronger tail dependency than you have over the middle points of the curve, the mean of the curve. So it does do a better job of maybe reflecting what one sees in the marketplace and then when things go bad, they go bad across the board.

MR. AHLGRIM: Improving the correlation, I think, is one of the biggest things we'd like to do.

MR. GORSKI: I guess the other thing is, in terms of analyzing the interest rate components of either your model or the Academy model, to remember that the Academy model was designed for evaluating interest rate risk. One of the important characteristics that was studied when the model was initially developed was yield curve aversion. And I'm not sure if you looked at both frequency and duration of varied yield curve aversions.

MR. AHLGRIM: We did something. The model allows it, but in terms of the prevalence of it, that could be something that can be easily checked.

MR. GORSKI: The point of calibration of the Academy interest rate model was calibrating to historical yield curve inversions.

MS. NICOLA BARRETT: In the diagram you had of all the interconnecting circles, real estate was off by itself with no connection to anything else. At the end, you said you can plug in your own proprietary real estate. Did you mean to have no connection?

MR. AHLGRIM: Actually, what we did is we checked for a relationship, especially between long-term interest rates and inflation. Those are two factors that we thought would be correlated and, in fact, based on the data that we have statistically they were not correlated. Part of it, we think, is the data that we actually got. Rather than transaction data, it's heavily influenced by appraisals, which tends to smooth out market fluctuations. So, and I mentioned it before, trying to find a publicly available data set for real estate transactions was very difficult. We can get it on residential transactions, but we wanted more commercial. We looked at it. We didn't see any statistical correlation, and we think that the big reason is because of the data set that we have.