RECORD, Volume 29, No. 1*

Washington, D.C. Spring Meeting,
May 29–May 30, 2003

Session 62TS
Building an Economic Scenario Generator

Track: Investment

Instructor: HAL WARREN PEDERSEN

Summary: Recent volatility in the capital markets has increased our focus on the risks involved in our businesses, and has underscored the importance of having tools to quantify and manage these risks. One of the most powerful tools for modeling risk is stochastic simulation. An economic scenario generator is among the vital components of risk management systems. Different generators can produce dramatically different characteristics for key macroeconomic and financial variables. This session focuses on the issues involved in building a generator. It covers the data set used to develop and calibrate the generator as well as a discussion of different approaches and their features. Implementation is addressed along with potential pitfalls. Examples of generator output are provided.

MR. HAL WARREN PEDERSEN: Let me begin with the part of an economic scenario generator that is probably foremost in most actuaries' thinking—interest rates. When any of us looks at different points on a yield curve, there's a remarkable amount of volatility. Another striking feature of interest rate data is that the trend has been down. Despite that, we still have continued to see varying shapes in yield curves. Just looking at shapes of yield curves doesn't give us enough information on what might be lurking around in the modeling process. One other thing to keep in mind through this entire discussion is the difference between modeling levels and modeling returns. Particularly in the interest rate world, getting realistic interest rate models may or may not yield appropriate returns. As you know, the returns are a fairly complicated function of what's going on in the yield curve.

When you look at a 10-year, coupon rate and a three-month T-bill rate, you can

see a downward trend, particularly in the more stable long rates. Obviously you can measure the slope by the difference. If it's negative, the absolute value of the difference is in the altitude. That's some basic interest rate data.
Recently, there's been a tremendous fluctuation in the long rate through relatively stable short rates. That has some important economic implications, but it also has caused us some real problems when we go to fit a number of models that you might use. That looks innocent enough, but when you really get into the data, there's always problems like that lurking around.

Another important issue has to do with the way equities behave. One of the things that we always are confronted with is whether or not our models are stationary. They may or may not be stationary. They may be stationary in unusual ways. The question that is obviously to people such as, personal investors, pension managers and everybody else is, whether or not the recent behavior of the S&P is unusual or just plain unpleasant. The next question would be whether or not the dynamics have changed. You can also get away from aggregating the data and look at what happens in individual stocks. My point here is that the question of stationary parameters and whether they need to be allowed to evolve in a certain way will be critical in the way your models are built and used. For example, a lot of clients don't want to see interest rates calibrated to their historical levels. They may want to see them stay low for some period of time. Other people may think that equities should be calibrated with low returns. Perhaps they're right. The only difficulty is we have very limited data to back that up. If you go to calibrate your model to something other than the history, then you're taking a guess as to what may happen. The question then remains of how you're going to structurally develop a model like that. Are you going to have a regime switching process or are you going to somehow tweak the parameters? Your guess is as good as mine. So as we go through this, keep these difficulties in mind, because they are paramount in applying these things in practice.

Let us look at the total return index for the S&P 500 on a logarithmic scale. In terms of its volatility, it doesn't look like it's changed materially over the period of time from January 1970 to January 2002. To really get into that, you'd have to look at the data on a much more detailed level. The one thing that is striking about this though is that we see what appears to be a prolonged bear market in the recent five years. Now, is this worse than 1973 and 1974? I don't know. Is this a sign of a change to come? My personal opinion from this is, I can't say. You may have a different feeling, but when we look at some individual stocks, there's different behavior in the last few years. Citigroup has been stalled. If you were to fit a straight line to the logarithmic returns, you would get a feel that it was going pretty good there for a while and now it's been pretty flat. Obviously the type of industry that you're dealing with is important.

Boise Cascade has paid dividends, but their stock hasn't done a whole lot. It appears that the Boise Cascade return process has a lot more fluctuations in it compared with Citigroup when measured in an appropriate fashion. I'm not sure that is significant, but when you look at the data, these are some of the things that would occur to you.

Wells Fargo is an unusual stock. If I contrast it with CitiGroup, it's not really the same behavior. Wells Fargo has, for the last six years, paid its dividends and stayed where it is. Of course, these need to be subjected to rigorous analysis, but these are the types of things you get an impression of when you're sitting down trying to build models.

Normally we don't model individual securities directly, but of course, individual securities make up aggregate indexes. Depending on the needs of a client, you might have to dig down further.

Another important issue has to do with the question of co-integration of economic variables. It's quite obvious that sometimes variables ought to move together. For example, you wouldn't expect inflation to be skyrocketing when the three-month interest rate was staying very low. However, we know that some variables ought to move together and some ought to move together in a much more restricted fashion than others.

I compared here two stocks that are both insurance brokerage firms—HRH and Marsh Mac. One is a much larger market cap than the other. Over the last three months at least, these guys have pretty much moved in lock step. Five years ago some Wall Street traders were actively engaged in "pairs trading" where they would select pairs of stocks that tended to move together, short the stock that was up and go long the stock that was down, because on average they tended to move back together. Over time the method broke down as more firms practiced it. For our purposes, the thing to keep in mind is that, at times, such co-integration effects have to be built into your models.

Just imagine you're modeling two securities like the case of HRH and Marsh Mac. They have to have some fluctuation between their behaviors because they're not identical, but you must do it in such a way that you don't have them drifting way far apart. If you examine the output from the scenario generator and look at a scatter plot of the returns, there should be some positive correlation, but you should still have outliers in any of these processes.

There's another issue about linking variables together and the more you try to link them, the more complex the models become. Unfortunately you tend to also get more parameters involved in the process. As a general rule of thumb, unless you think the linkage is fairly strong, you probably don't want to go to the trouble of embedding it. Unfortunately, in this kind of modeling, there's another problem that arises and that in that none of these relationships are ever that terribly strong. As I'm sure you're aware, if you try to explain the financial data, you always get very low R-squareds even if you do your regressions right. That's because if we could get high R-squareds from a lot of this, we can make money on it. You'd expect that most models we build wouldn't give real good explanatory power, but don't confuse explanatory power with the ability to simulate realistic distributions. We don't have to know where the market's going in order to simulate reasonable outcomes and assess risk. You don't want to bet your retirement or speculate a lot of money on it but for risk managers, like many of us are, it's fair enough to have a good feel for how these things move together and what their statistics look like.

Let's consider a classic type of regression where you look at equity returns versus changes in the T-bill rate. If you do this contemporaneously, you don't get very much correlation at all, but if you lag it three months, you get about -14%. It's significant, but there's some standard error. What does it tell us? It suggests that when the fed is dropping interest rates, stocks tend to perform better and vise versa. We have to be careful about how we interpret the causality as well. About all you can really say is that they tend to move together.

Now, what's my point here? Simply that when you build these models, you've got to grab some set of ideas together as to how you think things ought to move together and which ones you can ignore. Whether or not you can ignore this one probably depends to a large extent on whether or not equities are important in your portfolio or whatever problem you're trying to solve and whether or not bond returns matter to your present analysis. These are all problems that have to be addressed.

Now, going back to this issue of stationary parameters, we might ask ourselves if the volatility of returns on equity indexes stay constant. What about treasury bonds? While we all know the answer to that is "no," it doesn't stay constant. It is not unambiguous how you ought to measure it either. With most of the ways that you measure it, you wind up with certain outcomes. Again, the key point is that if you can identify a correct structural relationship and implement it, then you'll probably get good results in your model. If you implement a foolish or a poorly implemented structural relationship, which is my personal experience, because I've done lots of foolish implementations in the past, you get worse results and your model specification adds so much noise to the system that you're better off with a far more parsimonious model.

There is more than one way to measure the standard deviation in monthly returns. One way smoothes it more than some people would like, but the basic point is that the standard deviation of equities fluctuates a fair bit. As you would expect, medium-term treasuries have a lot less variability, but they both move. You might get away with using an average constant parameter for volatility if you're modeling these medium term treasury bonds. How you're getting at them, by the way, is another issue. Speaking in the most general case, for a process that doesn't move too much, a constant parameter may work. If you're really worried about is volatility, over some five-year period say, and what goes wrong during the middle of that interval doesn't matter, then you may get away with it. In contrast, if you're writing variable annuity business and you get hammered in between, having an inaccurate reflection of changing volatilities may be a real issue for you.

Let us discuss the changing of rolling correlations through time. The way it's interpreted, broadly speaking, is that you don't know which direction the returns of stocks and bonds will move together as you go through time and as of late, the direction has been negative. It may be that stocks have just done very poorly for so long that it's not surprising you have a negative correlation. I don't know. I haven't looked at it that closely. Again, you'll hear people quote the correlation between returns on bonds and stocks. It's nowhere near a constant number. Of course, if I look at it over an entire data horizon it's a constant number, but if I look at it over subsets, it varies a lot.

I think the foreign exchange market exemplifies the kind of problems you have in actually fitting data better than anything else. Foreign exchange is a very interesting topic, because we all have strong opinions about it. I'm Canadian, although I lived in the United States the last 13 years of my life up until this last January. You might think you can explain foreign exchange behavior based on relative prosperity, taxation and things like that. When you try to, for example, compare the behavior of the Canadian dollar against the U.S. dollar, you really can't get any explanatory power at all from such variables. Even though in the long run you'd hear economists say that things like purchasing power parity, interest rate differentials and other things like current account deficits ought to effect exchange rates, in the short run, they don't seem to matter. Furthermore, if you're always shocked into another equilibrium, do they ever matter? That's just something to consider when you think you can explain behavior in financial variables. Look at the data first, because it could very well be the case that none of these things explain anything.

If you test the U.S. dollar per Canadian dollar exchange rate using different unit root tests and so on, it really behaves like a trajectory of a random walk. That's about the best you can say. It's not a very satisfying thing to go and tell your boss, if they ask you to model an exchange rate, the best fit is a random walk. Nevertheless, it is worth knowing if it's true. It means your exposure to currency risk is something you have absolutely no ability to hedge except directly in the foreign exchange market. The most important thing you can do in such a case is to properly estimate the volatility. When I say hedge, of course you can go and buy contracts or swaps, but if you leave it alone, you're stuck with a random coin toss, in effect.

That's a look at some of the financial data that goes behind this analysis. Now, I'd like to talk about some of the tools that people use in this modeling process. I'd like you to just restrict your thinking for the moment to so called key financial variables. In a big economic scenario generator you're going to be trying to get at things like inflation, returns, interest rates, equities, perhaps mortgage-backeds, munis and maybe even have some tax interactions depending on the sophistication that your client needs. When you're starting out in this, the first thing you really ought to try and do is look at a couple of variables and see how you can use these tools to make them interrelate. For example, if you have stocks and bonds and you want to correlate them. Ask yourself how you're going to build a model that correlates the behavior of the two objects. Most models rely on a few fairly simple tools, because there are only so many things that you can do. You can certainly use more sophisticated mathematics, but there are not that many different avenues that you can take to link these objects up. There's some sort of a planning stage that goes on as well when you're assembling these models. I don't think it's too big an understatement to say that the most of the commonly used tools are, of course, stochastic differential equations and a variety of discrete-time econometric models. These models then get tweaked with additional things like stochastic volatility and autoregressive conditionally heteroscedastic (ARCH) and generalized autoregressive conditionally heteroscedastic (GARCH) modifications. They allow you to let the volatility move in non-constant ways in a continuous-time world instead of just having a stochastic differential equation. For example, for the stock price, you have some dynamics and then you might put in another equation to govern the volatility of the stock.

It's frequently much easier to write a model down than to estimate it. We can all very quickly build a spreadsheet that we could simulate a model with. It is very difficult to estimate the model correctly. At the end of the day, estimation is the key part of this whole process, and it can be unstable. Even though the tools you may use to write these models down are all within our grasp, the actual sophistication that goes into estimating and validating these models are very difficult.

Having said that though, a diverse set of modeling tools gives you a great deal of flexibility. If you're lucky, sometimes the nature of the problem you're dealing with is such that you might be able to get by with an estimate from a regression or some sort of a principal components analysis. Because regression an easier implementation, it's worth just having a quick look at what is involved in that. As you know, a vector autoregression is nothing more than a set of variables that have regressed on everybody else's legs. That process, of course, requires some know how regarding when you can do that and what adjustments may need to be made if it's not stationary. A related technique is the principal components analysis. Let's suppose we're only trying to model returns, and that might come up, if you're modeling variable annuities. You're not so worried about where interest rates have gone, you're more concerned perhaps with where the returns on the different mutual funds in the portfolio are going or how they're behaving. There are other problems involved in that, but this could be a situation where a simple first pass might work. You can do what amounts to in the ordinary least squares regression. You get a matrix for co-efficients and then you look at your residual co-variance matrix, usually take a Cholesky factorization and you can use normal disturbances to generate innovations to that and turn out your trajectories or returns. Now, depending on how the residuals look, which again can be in one looked at fairly easy with a graph, you may choose to use a t-distribution or perhaps a skewed distribution or some other means of controlling for any non-normality.

In my experience, these models can work well particularly over shorter periods of time. They may not be the best models. In fact, I'm sure they're not the best models to use for a 30-year analysis. The principal component approach (PCA) shares some of the same ideas. The main distinguishing feature of a PCA is that it lets you capture the most important variation of data using as few simulation variables as possible. Those of you who are familiar with this area might think of it as being analogous to the idea that we use a two-factor model to model default-free interest rates. Maybe we should use three or maybe you can get away with one. The fewer factors you have to simulate, the less dimensionality there is to the model. Ideally you get away with as few as necessary and still get the yield curves to twist and change appropriately. For example, I took corporate bond data, mortgage back data, SPX's, the total return index on the S&P 500 and some medium term treasury data. Some eigen values and some eigenvectors are going to pop out of that. The eigenvalues give you an indication of how much variance is explained by the first eigenvector, which is called the first principal component. If you do a little algebra you can put together the linear relationship for simulating these returns. I used only two factors to project the space of stock and treasury returns. We lost some of the dimensionality of the problem, but I reduced it to a two-factor model to explain the returns. I didn't take the statistical testing any further, but at least we're getting the negative correlations right. We're also seeing what appears to be a reasonable dispersion. For stress testing, I might like to see something a little stronger than a normal distribution in there, just to pop out in other directions and try to get the model to be a little more volatile. Of course I've chosen only two directions and that's fixed. As the volatility is increased or change a in the shape of the distribution occurs, you often wind up pushing the outliers in directions you don't want them to go. That's a limitation when you've cut down the number of simulation variables.

The projection on corporate returns and treasury returns are similar. The main difference is in default risk. There's not enough dispersion in the return difference between corporates and treasuries in this simple model. Default risk is a touchy subject, but without getting some way to make this fan out a little bit more, this model is probably inadequate. Of course, if default risk is low, then you'd expect to have a very tight band, but the way defaults behave in practice, as you may know, is that they tend to spike up over a period of time and then settle down and not do anything much at all. When they spike up you're obviously the most worried.

Let's take a look at equity returns. Equity returns are a real challenge in modeling because of the need to capture—not just the risk of bad things happening, particularly in the short term, but also changes in volatility. If you look at real data, then you can look for things like rolling horizon returns, and most investment houses and other data sources will assemble that. They'll take the history of, for example, the S&P 500 and compute rolling period returns and then order them from worst to best and produce what looks like a "horn-o-plenty." It gives you a feel for the historical range of the worst and best outcomes. It's not necessarily a validation tool directly, but you do want to see some bad things happen in the tail of your distribution. If you're using normal innovations, it can become a problem, because normal innovations are symmetric and that's going to mean that, as you increase the volatility, you're also wrenching up the upside potential. There's got to be some allowance to retain the downside potential without taking an overly optimistic look at what's going to happen on the upside.

We are going to compare this to a model that you may be familiar with developed by Mary Hardy and published in the *North American Actuarial Journal* (NAAJ). There's been quite a bit of other work done on regime switching, but Mary's paper is very clear and I think she certainly has a legitimate claim to laying down this particular type of switching model.

The only real thing you need to keep in mind on the difference between the geometric Brownian motion model and the regime-switching model is that the parameters of the log-normal returns are allowed to change through time in the regime switching model. We're only going to consider the two-regime model today, although you can calibrate it to more than that if you wish. Now, one strength and weakness with the regime-switching model is that the switching process is latent. It means that the economy turns bad or good based on something that's unobservable. In practice, it means that it's difficult to estimate, although there are a number of techniques you can do. Mary's paper provides parameters based on Markov Chain Monte Carlo (MCMC) estimation. At the same time, one of the other potential negatives with it is that we can't give any direct economic intuition to a latent variable. Having said that though, at the same time we don't often know why the economy turns bad or good. It's usually called the business cycle and usually ex-post blame is laid after the fact and it is probably appropriate to look at this as a latent process. I'll leave that for you to judge. I'll try to give you an indication of what this means and how it can be used.

I do want to emphasize that you can code up this type of model relatively easily and reproduce the results for the two regime model that Mary cites in her paper, but the actual estimation of this is very difficult and it would require a lot more work. Fortunately though, in the final analysis, what you really want is the simulation and not the estimation. So as long as you have some way to get a set of parameters that work, you can still use this tool even if you didn't know the estimation routines.

Mary also looks at the Toronto Stock Exchange Index (TSX). As you would expect, the returns are quite volatile and this 12-month rolling volatility is still fairly volatile too. There tends to be some periods where there's spiking. They're not very long, but they're there and they're pronounced. There was one in the 1975 area and one in 1987, 1988, 1998 and 1999. I guess the jury is still out on what we're going to see this year. In any case, one intuition on this data is that not only do we have changes in the volatility, they also tend to be accompanied by more negative returns.

One simple way to implement the models, if you know the parameters, is just to code up a little spreadsheet. I took the geometric compounded return over five years for the regime-switching model. Then, using the same innovations, so I can compare the behavior, I can go ahead and look at what happens to an ordinary geometric motion process. We thus have simulated returns for both models consisting of a histogram for returns over a five-year period. The most significant thing is that you get a lot more negative outcomes for the regime-switching model and you get about the same upside as you do in the geometric motion model, but because there's only 500 paths, there's a good bit of sampling variability. Depending on which sample you take, these numbers can change. Mary's paper actually gives a picture of the actual distribution and it looks much like you'd expect. By using a correctly parameterized regime-switching model, you can produce a returns density that captures some bad downside outcomes and, at the same time, not expand the upside potential for your model very much compared to a simpler model.

The regime switching distribution for this example is not more peaked than the geometric model. In fact, it's the opposite. It is possible that the distribution for stocks would have a fairly tight concentration around some low return number and a lot of thickness in the tails. In theory, you could have your actual data looking more peaked than the geometric motion model with thicker tails, but most of the models in practice tend to be flatter than the normal. Again, the key point here is to notice that the regime-switching type model will allow you to induce a different distribution of returns that lets you capture greater downside risk.

Just looking at returns can sometimes be a mistake. We also want to get a sense for what happens trajectory by trajectory. For the regime switching example, most of the time we are in state one where the mean return is about 1.25% and the standard deviation in that return is 3.5%. Once in a while we pop up to state two where the mean return is negative 1.85% and the standard deviation of the return is 7.5%. If you're in state one, you've got a positive return and a lower variance than you would under the ordinary geometric motion. Your variance is less and your return is greater. If you have the misfortune to wind up in state two, then you've got a negative return and much higher variance. That means that if I'm using the same innovations to trace out the trajectories, then we should expect that when both are in state one, all other things equal, the geometric motion process won't do as well.

Under both models we start with an initial investment of $10. Although they wind up with roughly the same amounts of money at the end because of the way the regime switching occurred for the path we're looking at, there were times when the regime switching asset significantly out performed the geometric motion and then significantly under performed it. Obviously that had to happen somehow if they wound up similar, but you can also see it from the difference of the slopes as we move through time. You want to pay attention to which state you are currently in because it determines the outcome of the period ahead. If you're in state one, then the parameters for that next period are governed by the state one parameters. With the second occurrence of regime two there's a negative draw because the standard deviation for the regime-switching model is much bigger. There's a much more negative force going down. That's the kind of thing the model is trying to capture. At the same time, if you have the misfortune to wind up in state two initially, then we're going to have plenty of downside risk even for short uses of this simulation. It's perhaps a reasonable model to use in order to capture some of the stock market risk a lot of people are worried about today.

Generally, it is important to look not only at return distributions, but at path-by-path behavior too. That's much harder to assess with a model, but don't forget that even if you get the distributions right, you sure don't want crazy inter-relationships going on in along individual trajectories. In short, you have to look at trajectory behavior as well as distributional behavior.

I think at this point it's worthwhile to talk about the concept of an economic simulation generator. For me, in practice, it means a way to consistently generate macro variables, particularly inflation, inflation indexes and the returns that I need to look at my insurance business. If you're in a business where some of your product lines are somehow connected to assets like price of crude oil, you want to capture those too if they're significant and if the effect can be actually measured and modeled.

In a generator you have some investment variables, such as treasuries and stocks. Corporates are very hard to model because of the default risks, but they're such an important asset class that any reasonable generator sooner or later needs to have them. Tax-exempt and mortgage-backed securities will make their way in there too. Insurance variables are included because many times inflation is such an important factor in the way your claims behave. GDP and unemployment can also be relevant and there could be other variables included too. Don't forget that inflation is somehow a very important link to everything that goes on in the bond market. If you need inflation for insurance lines, you've got look very carefully at how you're going to model those things together.

In many models you begin with the treasury term structure. Other models will start with inflation. I'll talk about one good model that does that. Generally you want to build your most difficult model at the top if you can. You normally cannot model all these things simultaneously, so you try to structure things in such a way that dependent items are grouped. For example, mortgage-backeds depend on interest rate behavior through prepayment functions and you try to make logical linkages that allow you to move down and get all the asset classes you need. Of course, sooner or later you probably need to make links to other economies too.

No sooner do you try to set all of this up then you start to have data problems. I'll let that speak for itself. I'm sure many of you would sympathize with that problem. Nobody will hand you the data you need. You think interest rates are interest rates until you actually ask about option adjusted amounts and coupon curves and you are suddenly unsure exactly what your data represents. Let's say what we're really after for the moment is some sort of a default-free treasury model and presumably you've got some type of idea on the type of the data that you're going to need. If it's a classical factor model, perhaps you're going to need some points on yield curve. If it's a market model, then perhaps you're going to need some assessment of volatility functions. You can get free data from the Federal Reserve Board. It's called the H15 data set for those of you who aren't familiar with it. Because bonds are issued at different points in time, sometimes these series are discontinued or they begin at odd times. With the termination of the 30-year bond there's now some other stuff in there called 25+ that allows you to make adjustments. If you know how the par-coupon quotes on the H15 data works, then you can use it to produce zero-coupon curves.

Some of the H15 yields are zero coupon yields and they're bankers discount quotes for bills, so you have to be careful how you make them consistent. Often times you need to get your hands on zero-coupon data. Probably the worst way to do it is to look at U.S. strips data. They never agree with the zero-coupon yields because they're not zero-coupons in the theoretical sense. You pay a premium for strips because they're a normally liquid marketable security and there's also some tax effects on them, although they can behave either way. That's not the same as what you get from stripping the H15 curve and any other U.S. treasury curve. Does it matter for your applications? Perhaps, but there's definitely a difference between strip yields and zero-coupon yields.

There are also problems over and above getting the data straight that apply to this kind of a model. And that has more to do with data changing on you and invalidating parts of your model. Of course, when that happens the problem is with your model and not with the data, but it sure can be frustrating. As I was saying at the beginning of this session, the recent behavior of interest rates has been such that we've really seen a lot of tension on more standardized interest rate models. If we ever get to very, very low interest rates, it will be very difficult to pick up anything meaningful when you've got long rates at less than 1%. As it stands, you can look at a lot of these models that were developed when interest rates didn't look anything like they do today. They stand up pretty well, but they do have their limitations.

Other things that matter in this too, particularly when you've had wide swings in interest rate levels or stock market behavior, are that if you try to estimate over very long periods of time, you're going to get too much tension on the model. If you've seen interest rates on the long end of the curve go from 8% to 3%, and you try to calibrate a model over that period of time, usually you get pretty bad results. The problem is that the model is mis-specified like any other model is, but almost any model is going to be really mis-specified with that big a change in interest rates. A lot of people will think that the more data they use to calibrate a model the better, but in practice, a lot less data usually gives better results. Up until recently, a rule of thumb might have been five to seven years of data, not 27 years of data.

What's quite unusual recently is that when you look at the three-month and five-year zero-coupon yields the three-month rate has basically been flat. It tends to be flat for long periods of time because of the Fed is fixing the rate and then it fluctuates a little bit in the secondary market. When you look at the fluctuation of the five-year rate, suddenly your model has to be able to accommodate a fixed three-month rate and get a bit of fluctuation in the five year. Now, many models can, but that's unusual. If you look back through time, you may have seen a little bit of that in the mid-1990s, but at the same time that the three-month rate is fixed, that five-year rate is hitting new lows too. Unless you work with this directly it's difficult to appreciate it, but we are certainly in some challenging times for model application.

I'll say a little bit more about the S&P 500. As I already commented on when we were talking about the regime-switching model, non-normality constantly becomes an issue. You might argue that there's some non-normality in interest rate data too. In fact, it's pretty clear that there is, but when you get down and look at the innovations, you may be able to use more normal-type processes if you've transformed them correctly to begin with. The only thing to keep in mind is that you shouldn't assume your data is going to look normal.

Another solution that doesn't require the same sophistication as a regime-switching model would be to use an adaptation. You splice into the tail of your innovations, something like an extreme value distribution or even an exponential distribution. When you're splicing it, you have to make sure you preserve the overall mass or you're doing something crazy. You can thus alter upside and downside potential in your returns. This isn't very far from a value-risk type analysis if you don't have any dynamics for the model, but this can be made to have an innovation term that's not normal and still produces the same general behavior.

You can also do total return data for government bonds, corporate bonds and all that other good stuff. Unfortunately, the dividend yield series for the S&P 500 is hard to come by now as it has been discontinued by standard vendors. In practice, if you have to model equities, you're going to run into a data problem because dividend yields will be needed. Sometimes, all we could get was what amounts to an implied dividend by taking the difference between the total return index and the price index. It looked pretty unstable. You can take a 12-month rolling average, and it calms down okay, but now it's pretty smooth. This is a good example of a real world problem where suddenly the data you wanted to calibrate with your model disappears.

When it comes to looking at total return indexes on bonds, there are many ways to get that. You'll have a devil of a time getting individual corporate bond prices, because that's highly proprietary. You can get total return indexes from companies like Merrill Lynch. We've generally found them to be very good. One of the more striking things about the corporate market is that, at times, defaults start to fluctuate and returns on lower quality corporates can become very volatile. A lot of people have been piling money into high yield lately. When credit quality deteriorates, that sometimes leads to extremely bad outcomes.

We looked at historical default rates from Moody's reports aggregate statistics. We looked at all corporate-bond default rates and speculative grade corps. As you know, speculative grades are those below BBB, so they aggregate them all together. On the junk front you're going to get a lot higher rates. You don't see stable and level default rates. What you see instead are spikes. Lately it appears that the spike has subsided and, in fact, if you were to plot recent data you'd see that it's definitely backed down sharply. If you're going to model corporate bonds, you surely can't just take an average default rate and hope for the best, because trajectory-by-trajectory, you're not going to see anything like what happens in the real world. For such a simple model, you may make asset allocation recommendations in certain asset classes that work okay on average, but path-by-path they're disastrous. It's not good enough to get the average distributions right. We absolutely must look at what occurs on a trajectory-by-trajectory basis as well.

High yield returns were a nasty looking series, although I suppose it depends when you bought it. An example of a notorious bond is WorldCom's bond. Our model has ratings modeled in it. Ratings are not anticipatory things though, especially in practice. They tend to happen after the fact. I have a great deal of respect for all of the bond-rating agencies. They are talented people. In fact, it's an area that actuarial training is well suited to. However, in this instance there are definitely some practical problems. Note here where your last quote is at 15 and it's just been downgraded. You wanted that information before the price went down if at all possible.

We looked at causality in the ratings changes. We don't know which way it runs, but it could be that the rating downgrade is merely confirmation that the bond isn't worth much any more. What is truly shocking is how rapidly the price of that bond comes down and that's absolutely typical of what happens to debt when it gets downgraded into the junk classes.

We also took a look at historical statistics on returns by rating class taken out of Merrill Lynch's indexes. We had a clear indication that you're far and away worst off in junk stuff, because your average monthly return is so low and it had the highest variance to an efficient outcome. Having said that though, we looked at averages. They're not what's going to happen on any given trajectory, but the average behavior of these suggests to me that you're best off probably in an A or AA bond. AAA is a little bit anomalous in some ways because there's so few of them now.

There have been some good studies done at the SOA about junk bonds and corporate private placements. You have to be very careful with what the structure of these is and that's just looking at history. Imagine what it's like to try to model the returns. Another issue is getting transitions right. You definitely need to model transitions in this whole process too. One of the hard things is getting the defaults to look right and there's a lot of mass on the main diagonal with either matrix. The real aggregate data has plenty of mass on its main diagonal. Whatever your mechanism is for defaulting, transitioning has to be done in such a way that those masses are preserved. You also have to be able to recalibrate the model as the market conditions change.

When it comes to validating some of these models, you can look at what the maximum returns were over different periods of time and over different asset classes. Validating gives you an idea of where the relative returns should sit if your model is calibrated to something that looks like history. It's in the short run that you're going to see the greatest variability and, as you hold these different asset classes longer, they'll converge to some number. If you're looking of course, at actual historical numbers, it converges to the long run average.

If you're going to try to build a simulator, a platform is going to be important when you're actually simulating this for real. An Excel spreadsheet is a good place to play with stuff, but it's not a very good place to run a simulation. I don't know what programming language you might find suitable, but that's an important part of your decision process. The value and the limitations of a generator also need to be considered. First and foremost are things like enterprise-wide risk management and a caution that these generators are not predictors. They're designed to assess business performance across a range of plausible scenarios. No generator is going to produce all the dooms-day scenarios. There are still scenarios that can be benchmarked against history. If somebody comes to you today and says I want low stock returns, the honest to goodness truth is you're taking a guess at what stocks are going to behave like. In a regime-switching world, you might crank up one of those transition probabilities or tweak one of the mean or standard deviation parameters, but you're guessing. You're not estimating. You can simulate it out and see what it does in the long run and look at the trajectories and that's all good, but the real problem you have your hands is that you only got real data and subjective opinions and there's got to be some mixture as to how to make them work. Depending on what the biggest fear in your business line, or your company is, you have to make some judgments as to how to incorporate these bad outcomes.

I didn't say much about arbitrage-free modeling and I'm not going to, but it's possible in the contexts of working in many arbitrage-free models. At some point, the more assets you get, the harder it becomes to maintain that. In our model we use at dynamic financial analysis (DFA), we've made some compromises and some subsections in our model are not strictly arbitrage-free, but they make economic sense. They work well and that's not much different from what is done in any other firm.

The output I will show you is from our system at DFA. We have bonds, treasuries and mortgage-backeds—all of the usual types of variables. They're arranged, as we've already discussed, in somewhat of a cascade structure. Treasuries are the key driving variable. Modeling artificial economies require a trade off. First of all, joint distributions definitely have to look like something we could associate with history and they better have some correlations that we think make sense too. Arbitrage-free is towards the bottom of the list for real applications. If possible, you'd like it to be arbitrage-free, but if we can't save that, then my personal opinion is that it's better to work with a model that fits the data and produces reasonable simulations. Having said that, don't mix arbitrage-free with option pricing. If you're using this model for option pricing, that's a very different set of issues. For the moment and in this discussion, I'm only speaking about using these models for simulating economies. When you move to the arbitrage-free pricing world, there's a whole different set of problems that are relevant. You can't price options unless you have some sort of an arbitrage-free type model or at least something you can relate to that context. Don't mix the two up.

The treasury model has a structure that looks somewhat like a two-factor affine model. Our stock and dividend yield model is coupled to that and it has some additional sources of risk. The corporate model is something I've spent a lot of time with. That is linked up to the treasury model and the stock model. It's also related to default rates. Quite honestly, it's a fairly complicated model. The mortgage-backed model, as you probably know if you've ever tried to work with this, is also a difficult thing to get right. Not only do you have these cash flows that are prepaying on you, you also have an option value embedded in the prepayments that has to be priced in order to get the appropriate yields on these things. Both of them depend on the treasury term structure, which you can model that a variety of ways, but there needs to be a linkage. If we have a multi-country model, the first linkage is made through treasury and equity markets.

Here is an example of BBB yields for our corporate-bond model. What happens in practice is that we get a smattering of yields on different points on the curve. For any given trajectory, each bond will be modeled as an individual security and we'll know what its yield is. At any given point in time, you can take this bunch of individual securities and assemble an average yield curve, depending on which fitting criteria you want to use. Most of our clients are interested in corporate returns and they're holding corporate bonds in order to generate cash or pay liabilities, so it's important for us to track what has happened to their bonds and portfolios when we move through time. To get a feel for where the market is, you may want to be able to produce a yield curve. Beware if you don't like having multiple securities, because you're going to have a corporate bond model where all bonds at any rating will have the same price. In the treasury world, all treasuries are interchangeable securities if they have the same maturity and coupon. In the corporate world, every security is different, because every company's credit rating or every company's prospect is different. They can be very similar at times, but there's no such thing as a non-ambiguous corporate-yield curve at any point in time. This is no different than what actually happens in practice.

With a generated of model high-yield data you should expect to see some outliers, because not only do these volatilities of the yields increase as the rating category drops, but the rating classes will overlap at the same time. You will see actual BBB bonds that have a lower yield than a single A for the same maturity.

You're also going to want to know what a sample simulated transition matrix looks like as a feature of the corporate model.

Our system has a built in estimation routine. Many companies do not. This, I believe, is one of the biggest accomplishments of our system above some of our competitors in that we can actually estimate on the run from our database. We've coded all the necessary algorithms as part of the system. A lot of companies have whole departments that play with perimeters to try to produce estimates. This thing we use actually estimates in real time. There are quite a number of parameters to be estimated for one of these treasury equity complexes but we are able to handle it.

There are quite a few perimeters for the corporate-bond model too. One way that we validate things in the end is we look at the average mean variance frontiers. For a mean-variance frontier for our system, we look at all the returns and variances along each path in the simulation. If you're running an economic simulation engine, then one of the things you probably want to do is compare the returns on your simulated asset classes with the returns over some period of time that you think is representative of history. This is an important part of a validation procedure.

It is lot of work to set up reports like the ones from our system, but if you get into trying to build these generators, don't forget the importance of validating what you do. We all talk about it in our work, but once these tools were built, we had a great way to be able to compare what a change in some parameters does to the system output. Without such diagnostics you really don't know where things are set, especially when you get a whole lot of different models interacting.

There are some real interesting issues that go into setting these models up, but I'd start with an interest rate model or a stock model, learn about the distribution and then begin to try to tie all the pieces together.

FROM THE FLOOR: How did you determine the parameters when using the regime-switch model for the probability of switching from regime one to regime two, or the probability of staying in the same regime once it's in there? How do you assess those parameters? You said that if the long-term interest rates are very low, which is very true, the interest rate model wouldn't be able to capture the simulation for the future. It wouldn't be able to work very well. Are there any suggestions to get around of that problem?

MR. PEDERSEN: I took those parameters directly from Mary Hardy's paper in order to use that illustration, but if you look up her paper, which is in NAAJ and available on the SOA Web site, she estimates using the Markov Chain Monte Carlo method. Another way to estimate the model is called the expectation maximization (EM) algorithm, which involves classical maximum likelihood. In any rate, the answer to your question would be to begin by looking at what Mary Hardy describes, because she does outline her algorithm. If that one isn't satisfactory, then I would try an EM-type algorithm.