RECORD, Volume 26, No. 3*

Chicago Annual Meeting
October 15–18, 2000

Session 27PD
Interest Scenarios

Track: Education and Research/Investment

Moderator: FAYE ALBERT
Panelists: SARAH L.M. CHRISTIANSEN
VLADIMIR S. LADYZHETS†
JOHN D. MARCSIK
Recorder: FAYE ALBERT

Summary: Actuaries are increasingly incorporating interest-rate generators into their analysis. Insurance and finance practitioners discuss purposes and approaches for creating interest-rate scenarios.

Ms. Faye Albert: We have a distinguished panel of experts on interest-rate generators and their uses. All of our panelists are Ph.D.s.

John Marcsik is with Nationwide Financial and works in asset/liability management (ALM). He will be giving us an overview of interest-rate generators, different models you might choose, and what the advantages and disadvantages of the various models are.

Vladimir Ladyzhets is working at MarketSwitch. He's from Russia and has been interested in financial models ever since he got to the U.S. He'll be expanding on creation of the parameters for a model and ways you can determine what your model looks like in more detail.

Sarah Christiansen, our last speaker, is with the Principal Financial Group. She'll be telling us how insurance companies have used ALM and interest-rate models, and what regulatory concerns have precipitated that type of analysis.

Dr. John D. Marcsik: Today I'm going to talk about ALM factor models, as they're generally called in the literature. The advertised session is called "Interest Scenarios," but most of the ideas that we will cover today apply just as well to any market model, whatever index you're trying to model—S&P 500, currency exchange, or anything of that nature.

What is the purpose of using these models—what are we trying to accomplish? What are the features the models have and why might they have them? What is the scope of such a model?

---

†Dr. Ladyzhets, not a member of the sponsoring organizations, is a Senior Research Analyst with MarketSwitch Corp. In Farmington, CT.

There are no financial market models that are perfect for every possible use; each of them have an appropriate scope and appropriate interpretations. The real benefit of these tools is in increasing your understanding about the way you're running your business—in particular, the way you're running your ALM. I'll go through an example of the thinking process on how you'd create such a model. In the end, you have to make some trade-offs because markets are enormous things—perhaps millions of participants and tens of thousands of securities. This model is not going to try to capture every nuance of the behavior of millions of people. Try to pick out what's most important.

The purpose in my mind for creating these models is to do ALM analysis, i.e., analyzing different strategies, different pricing, and different ways of dealing with things and analyzing the risk/return profile. The output of the models is usually some sort of distribution so you're measuring some sort of risk and return trade-offs.

This kind of modeling uses different disciplines. The four mentioned are economics, finance, statistics, and mathematics; it's a nice blend of all of those. Economics. You get supply, demand, and how markets work. Finance. Specifically, what types of securities (in this case insurance company) would be suitable investments? Statistics. All these models are statistical models, which means even after you specify a mathematical form, statistical fitting of parameters is necessary and all the results are in statistics. Of course, to make sure these models are internally consistent you've probably heard the term a million times and no arbitrage. I prefer just to say internally consistent models so that you've also modeled the risk you're taking as an investor to get that reward. And, of course, mathematics is how you solve for the security prices.

It would be nice to make very complicated models and hit every nuance of every yield curve in history, but chances are if you created such a monstrosity you would never solve for a yield curve. I'm trying to build a stochastic model of yield curves. I have to be able to solve for the actual yield curves that are implied by the model.

We try to model a market, say an interest-rate market, with few state variables. We don't want to model every Treasury that's on the market that's trading, or every corporate bond (that would be in the thousands). What we want to do is categorize them: different maturities, different credit qualities, possibly different options embedded in them, and create a few state variables. You call these factor models. All of these prices and all of these securities are functions of those state variables. It's very rare that you can come up with a Black-Scholes formula where it's a closed formula. Usually the model is some kind of numerical algorithm. For this purpose I prefer to use equilibrium models. Vladimir will get into four different types of interest-rate generators. These are equilibrium models.

If there's some short-term thing going on in the market because of some supply and demand, but it's short-lived or cyclical, then maybe it's not modeled very precisely in these types of models. When you solve for a yield curve in these models, you never hit the price of every bond sold. You try to get as close as possible with a few number of variables. Try to capture the essence of the pricing.

I use both what people call realistic and risk-neutral models. You want to make sure that they're consistent; that there's a realistic probably that you're using, and the risk-neutral pricing to solve for the yield curves you're using, are compatible so that you don't have any built-in rewards without taking some sort of risks. Here the risks are on the state variables that we're modeling. For each state variable, you can think there's some risk premium that goes into the pricing of securities whose prices depend on those variables. The key there is to make sure you're consistent. If you have a realistic model and you specify functions for risk premium, there are mathematically defined risk-neutral prices, so you want to be consistent in that way.

For ALM, a derivative trader might use different products that have different time horizons, which are probably longer than an interest-rate model. We're probably talking about significantly longer time horizons.

You can think of this type of stochastic modeling as an evolution. Think back to before everybody got a PC. ALM was matching duration because that was a calculation you could do without a computer. If you didn't have options embedded in there, and even if you did, you probably didn't have any choice but to do it that way. With duration you hit your model with an instantaneous change in your variables, whether they're yields or other state variables, and you look at the way your assets and your liabilities values change with that instantaneous hit. There are lots of weaknesses to that. First of all, duration usually is defined with parallel yield curve shifts, so even if you immunize yourself to parallel yield curve shifts, the yield curve can tilt, flex, change shape, or invert and that type of analysis doesn't pick up those risks. Duration calculation doesn't pick up large changes. It's an infinitesimal calculation so if interest rates jump 100 basis points, you might miss that.

It also doesn't take into account how the values of the assets and liabilities are changing with respect to time. If you take the mathematical derivative with respect to time, you don't necessarily have that matched. Simple calculations will miss any sort of optionality that is embedded in either your assets or your liabilities.

A stochastic model tries to do the same thing, but it looks at many changes, not just instantaneous infinitesimal parallel yield curve shifts, but many changes over different time horizons. It creates statistics for all these things. And it prices securities whose prices depend on those variables consistent with the statistics and a risk premium function that you specify.

From my perspective, a lot of the things that actuaries might do should include ALM analysis. Certainly investment policy for an in-force block of business, where you have a portfolio manager who is investing assets, is an obvious place for ALM. Pricing. These types of interest-rate scenarios can be a part of pricing whether it's securities or liabilities. Product design. Actuaries love to put options in all their products. The only way to fully analyze those types of things is with a full-blown stochastic model. Capital efficiency and finally hedging. In hedging there's another requirement hoisted on your model; it has to be tied to real markets. It's not tied to real markets the way a trading model would be, but you want to optimize your model so that it's very close, as close as you can be, to actual market prices. So if you're devising a strategy to hedge things, you don't want your model saying that certain puts or caps or whatever costs X when the market really charges you X plus 30%. You will make very different decisions if you price options in a way that the market doesn't.

Here are some desirable features. I've already mentioned some. You want to be internally consistent and theoretically sound. Most of that has to do with arbitrage-free considerations. That's simple. Of course, given two models, all things being equal, simpler is better. Remember, the purpose is to increase your understanding. If the model is too complicated, you won't increase your understanding at all; you'll just be confused. Modularity is a good feature, so if you have a model that does Treasury rates and a model that does corporate bond rates, it's possible that the part of the model that prices the credit risk can work equally well with two different interest-rate models. That would be preferable. That way if some day you change your interest-rate model, you don't necessarily have to change your credit-pricing model. That's what I mean by modularity.

Long-term. We're talking about insurance companies mostly and long-term liabilities, so these models should be long-term conscious. Mostly that's in how you calculate the statistics the model is going to output. You might look at historical prices over long periods of history and historical yields over long periods of history.

Computationally tractable. For the most part what I'm referring to here is prices. So you have risk-neutral prices. How are you going to calculate them? There's some trade-off there between statistical fit and being able to price quickly.

Log-normal is only the first approximation; it can be misleading. In my experience log-normals, except for very short time frames, fit very poorly. They're wrong in skewness; they're wrong on tail risk. It can be misleading. Even though you might think of it as a first approximation, you have your mean and your standard deviations. When you're talking about one-sided risks, they can be very misleading.

In comparison with a trading model, a trading model tries to take a snapshot of a market and probably every asset in that market and fine-tune parameters—usually time-dependent parameters—to reprice every security exactly or very close to exactly, which means that if in the middle of the day some large mutual fund starts buying something like crazy and actually moves the market, we would say this is not an equilibrium but one buyer is temporarily skewing the prices. A trader needs a model that picks up on these things. Perhaps that's an opportunity for a trader; for an ALM manager, that's not an opportunity. They're not investment managers. They're doing more long-term, what I call equilibrium-type analysis.

Some of the things you might want to model, for instance, are Treasury curves. If you model for credit spreads based on Treasury curves and you know this year

Treasury curves have been kind of strange, possibly basing something off swap curves makes more sense. These types of models tend to work better when you're looking at very liquid markets, so the less liquid the 30-year Treasury market becomes, the harder it is to model it.

I thought I'd go through the thought process and how you would use a corporate bond model. Again, you don't want to price every bond exactly. You want to categorize them. Ratings, industries, and maturity lengths are obviously related to Treasury bond prices and they are not perfectly correlated. So for this I'm going to assume we have something that prices Treasuries and uses either a short rate factor for pricing Treasuries or the standard textbook type way of pricing Treasuries. You might add a few more state variables such as a default rate or a transition rate of credit ratings. And you need to specify some kind of risk premium that the market would charge for these types of risks. You probably also want to consider liquidity since corporate bonds are less liquid than Treasuries.

There are two types of corporate bond models that I know of. Then you would try to fundamentally model the company that is backing it so you model the assets of the company. When they get too small, it is bankruptcy and model it that way. Or I like to prefer a reduced Faure model because here we are not modeling any one company. We're trying to model in general corporate bonds and how they're doing, so for a reduced Faure model you don't need to know anything specific about the company, just probability of default and what the recovery rates would be. Those models also come with very good nearly closed forms for calculating yield curves, which helps.

Modularity. I think I mentioned that already. If possible, you'd want to make sure that you could replace your Treasury curve model without undue extra work in changing your credit model. There are lots of trade-offs to consider. You can add more factors. Of course, it becomes more complicated and harder to understand with more calculations. On the other hand, you get better statistical fits and model more risks. If they're more modular, you can reuse them a lot. There's less to intertwine, so they're easier to understand. On the other hand, if you're missing some true statistical relationships or some structural relationships, then that's obviously a trade-off.

It's internally consistent to price these things in a risk-neutral way. You can understand that. You get a reward for taking a risk. However, it makes the computations much harder. With that I'll turn it over to Vladimir.

Dr. Vladimir S. Ladyzhets: As you can see from the title of my presentation, "Term Structure Models: Which One Is Right for You?," I would like to talk about different types of term structure models. And if you asked, can I do my presentation in one sentence, the answer is yes. The sentence is, "There is no such thing as a term structure model which can work for all purposes you might need in your work." I would like to elaborate from that and show you what kind of term structure models you can choose from for different applications. This is intended to be educational, but I hope that you will also find it practical. So let's start with a description of term structure models that are available.

If you talk about different term structure models, you basically have two choices: you might have an arbitrage-free model or an equilibrium model. Both types were mentioned by John. And for each type of model you might have two parameterizations: realistic and risk neutral. In total you have four options, which are presented below.

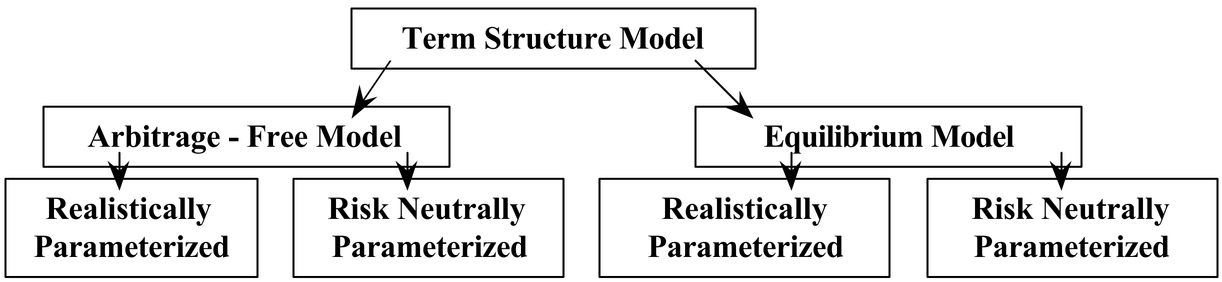

Let's talk about all the four options or combinations. Let's start with an arbitrage-free model. What is the purpose of this model? As John already mentioned in his presentation, the purpose of an arbitrage-free model is to be in sync with the market. You don't want to have the price above or below the market price, because as soon as we have this you provide an arbitrage opportunity and you would like to avoid this. So what does an arbitrage-free model do?

An arbitrage-free model takes certain market prices as given and adjusts the model parameters in order to fit the prices exactly. If you have 100,000 of market prices, you choose some of them as the basis. After that we fit the model in such manner that all other prices match what you see on the market.

An arbitrage-free model assumes some computationally convenient random process underlying the yield curve. Add time-dependent constants to the drift and volatility until all market prices are matched. Something happens to the prices, somebody starts selling or buying, what you would like to do is synchronize your prices with this move. So that's why you need time-dependent parameters or time-dependent variables.

An arbitrage-free model requires at least one parameter for every market price used as an input to the model. This is the price you pay. If you would like to have more securities being priced by your model, then you should have more parameters in your model. This means you have more headaches when you are dealing with the model, but you are more accurate. It is a trade-off and there's no way out.

An arbitrage-free model looks like an excellent type of model, but there is a dangerous aspect to using this model. It does not actually attempt to emulate the dynamic of interest-rate term structure. An arbitrage-free model produces reasonable current prices for a particular class of assets by interpolating among existing prices. Basically the prices you are trying to match are like points in space and by using an arbitrage-free model you draw the hypersurface that goes through all of the points. It is important to be aware that by doing this you are not trying to predict what's going to happen with these prices.

- Ho & Lee Model (1986): $dr = q(t)dt + s\,dz$
- Hull & White Model (1990): $dr = k[q(t) - r] \cdot dt + s\,dz$
- Heath & Jarrow & Morton Model (1992)**:**
  $dF(t,T) = m(t,T)dt + F(t,T) \cdot \sum_j s_j(t,T) \cdot dz_j$

r is a short-term rate; *F* is a forward rate; *dz* is a Wiener process.

I would like to direct your attention to the fact that all the coefficients in the arbitrage-free models are functions of time *t*. Among the three models that are presented above, the Heath & Jarrow & Morton Model is probably the most famous and most generally used. The principal difference between the first two models and the third is that the Ho & Lee Model and the Hull & White Model are written as equations on the short-term rate, while the Heath & Jarrow & Morton Model describes the dynamic of the forward rate.

All three models are classical and rather old. New, more advanced models have recently come along. Regime-switching models as well as models that combine Wiener processes with Poison ones and are especially tailored to deal with large changes in the interest rates.

Now if you are talking about equilibrium models, these models have been developed to deal with a situation which is quite different from one when you want to apply an arbitrage-free model. You are not intending to match all of the prices. You are trying to catch the tendencies and the movements of term structure.

An equilibrium model employs a statistical approach to the price observed on the market. Basically, when dealing with equilibrium models you are saying that an observed market price is a real price plus some stochastic error, so we do not expect that the equilibrium model will match all the market prices as they are observed. This is the major difference between an equilibrium model and an arbitrage-free model: an arbitrage-free model takes prices as given while an equilibrium model estimates prices.

As you remember, when dealing with an arbitrage-free model you need to handle a lot of parameters. An equilibrium model uses a small set of step variables, which are fundamental components of the interest-rate process, to describe the term structure. Compared with an arbitrage-free model you are losing accuracy, but the model becomes easier to handle, so it is easier to understand what is going on and easier to interpret your results. An equilibrium model does not contain time-dependent parameters; instead it contains only a small number of statistically estimated parameters drawn from the historical time series of the yield curve.

In order to calibrate an equilibrium model you are analyzing the available historical data, let's say for 10, 15, 20, or 30 years, depending on your time horizon, and you try to extract some statistics from the data. Usually people deal with such statistics as volatility of yields for Treasury securities with different maturities, spreads, and correlations between different positions on the yield curves and so forth. When calibrating an equilibrium model you are trying to find coefficients for your equations that would facilitate a good match between the statistics in the stochastic scenarios produced by your model and the similar statistics that were observed over some historical period.

Below are some examples of equilibrium models:

- Vasicek Model (1977) |Ornstein-Uhlenbeck Process (1930):

$$dr = k \cdot (q - r) \cdot dt + s \cdot dz$$

- Cox & Ingersoll & Ross Model (1985) | Feller Process:

$$dr = k[q - r] \cdot dt + s\sqrt{r}dz$$

- Longstaff & Schwartz Model (1992)**:**

$$dr = k_r [q_r - r] \cdot dt + s_r \sqrt{r}dz$$

$$du = k_u [q_u - u] \cdot dt + s_u \sqrt{u}dz$$

r is a short-term rate; *u* is a long-term rate; *dz* is a Wiener process.

Among the three models above, the Vasicek model is a classical one, which is probably the best known. As you can see, it's based on the Ornstein-Uhlenbeck Process, which was investigated in 1930. The next model that is also classic is the CIR Model, which is based on the Feller process. The third model, the Longstaff & Schwartz Model, was developed in 1992 and is the classical and simplest form of interest-rate models. Other, rather complex equilibrium models are available now.

Now let us talk about different model parameterizations. We start with the risk-neutral parameterization. It is based in the Girsanov Theorem, which is a very fundamental mathematical or statistical fact that underlies the risk-neutral parameterization. This theorem was proved in the 1960s.

The idea is that if you have a stochastic process *X(t)* that is defined on the probability space *(W,F,P)* and meeting some conditions, then there is a unique probability measure *Q* such that:

- the measures *P* and *Q* are equivalent;
- the stochastic process *X(t)* is a martingale, i.e. zero-drift stochastic process on the probability space *(W,F,Q)*.

We know that any price and any interest rate is a stochastic process defined on a probability space, which specifies the way the probability of different events is measured. The Girsanov Theorem says that you can replace the measure *P* by the equivalent measure *Q* in such a way that when measured by *Q* the process becomes a martingale. Roughly speaking, the measure equivalency means that if some event is impossible in one measure, it's also impossible in another measure.

What does it mean to replace the realistic measure by a risk-neutral martingale measure? In convenient terms, this neutral parameterization brings us from the real world where an investor expects (frequently in vain) that the extra risk will be compensated by extra return, to an imaginary world where there is no extra compensation for the extra risk taken. So the question arises: Why do we use the risk-neutral parameterization?

I would like to use a very simple example in order to explain the risk-neutral parameterization in non-confusing terms. Let

- r_t be the short-term rate at time *t;*
- *D(t,T)* be the price at *t* of a discount bond paying \$1 at time *T;*
- *S(t,T)* be the spot rate at *t* for the term *T-t;*
- *f(T-t)* be the term premium required by the investor at time *t* for the term *T-t.*

The price of the discount bond *D(t,T)* can be expressed through r_t, *S(t,T),* and *f(T-t)* by the formula:

$$D(t,T) = \exp\{-S(t,T)\cdot(T-t)\} = \exp\{-f(T-t)\}\cdot E\left\{\exp\left[-\int_t^T r_s ds\right]\right\} \qquad (1)$$

The equation (1) answers the question: How much should you pay at time *t* in order to be paid back \$1 at time *T*? In analyzing this, there are two factors that define the price: the risk-free, short-term r_t and the term premium *f(T-t)*. If there were no risk because of the changes in the interest rates between the purchase time *t* and payoff time *T*, then the price would be just:

$$D(t,T) = E\left\{\exp\left[-\int_t^T r_s ds\right]\right\}$$

But since the buyer is exposed to the risk of interest-rate change, he or she requires an additional price discount equal to: $\exp\{-f(T-t)\}$.

What is important for us in equation (1) is that two discount processes, one because of the risk-free, short-term rate and the other because of the term premium, are presented separately.

Let us modify the risk-free, short-term rate r_t to the so-called risk-adjusted, short-term rate by the formula:

$$\mathbf{r}_s = r_s + f(s-t) + f'(s-t)\cdot(s-t)\ . \qquad (2)$$

I just would like to note that the formula (2) presents a very simple form of the Girsanov Theorem. It is easy to prove that the price *D(t,T)* can be calculated by the formula:

$$D(t,T) = \exp\{-S(t,T)\cdot(T-t)\} = E\left\{\exp\left[-\int_t^T r_s\,ds\right]\right\}. \qquad (3)$$

Now we can see what is the idea behind moving from the risk-free rate r_t to the risk-adjusted rate r_t. We take the market risk and embed it in the short-term rate.

When comparing the equations (1) and (3) we can see that valuing assets using risk unadjusted r_t requires the more complicated discounting procedure which applies additional discount factors to the short-term rate paths to compensate for market risk. Under the risk-neutral parameterization the term premium, whatever their values, that exist in the marketplace are embedded in the interest-rate process itself. So the expected discounted value of a cash flow at the risk-adjusted short-term rate is equal to the discounted value of the cash flow at the spot rate. So you make your computation easier.

In realistic parameterization, you distinguish your market risk (term premium) from the risk-free, short-term interest rates. When generating interest-rate scenarios under the realistic parameterization, you create the realistic environment where more risk taken assumes more premium given.

Now that we have described all the types and parameterizations of term structure models, we can answer the question: How can one choose the correct model? We can combine all possible options into the matrix, which is shown below:

Classification Parameterization	Risk-Neutral Parameterization	Realistic Parameterization
Arbitrage-Free	Current pricing, where market prices are reliable.	Unusable, since term premium cannot be reliably estimated.
Equilibrium	Current pricing, where market prices are unreliable or unavailable. Horizon pricing.	Stress testing. Reserve and asset-adequacy testing.

Looking at this matrix we can see which one is applicable for a particular purpose. Let us take the top left cell: arbitrage-free model with risk-neutral parameterization. It's perfect for current pricing where market prices are reliable because you are matching the current market environment. If you take an equilibrium model with risk-neutral parameterization, which is the bottom left cell, you can use it for current pricing where market prices are unreliable or unavailable, or you can use it for horizon pricing. An arbitrage-free model with realistic parameterization, the matrix top right cell, appears unusable since term premium cannot be reliably estimated. And finally the matrix bottom right cell, an equilibrium model with realistic parameterization, is suitable if you are interested in stress testing and for reserve and asset-adequacy testing. Basically that's what you are going to need to conduct ALM.

I would like to illustrate four combinations of models and their parameterizations we just discussed by modifying the classical CIR model. We can call them four faces of the CIR Model. Let us start with the risk-neutral and arbitrage-free form of the CIR model:

$$dr = k(t) \cdot [q(t) - r] \cdot dt + s(t)\sqrt{r}\,dz \;, \quad r_0 = r(0).$$

As you can see, all the coefficients are dependent on time. And the question is, how can you calibrate $k(t)$ and $s(t)$? In order to do that you match cap or option prices. As for r_0, which is your starting point, and $q(t)$, which is the mean reversion parameter, you calibrate them by matching to bond prices.

The second face is the realistically parameterized arbitrage-free CIR version:

$$dr = k(t) \cdot [q(t) - l(r,t) - r] \cdot dt + s(t)\sqrt{r}\,dz \;, \quad r_0 = r(0).$$

In this case, the parameter $l(r,t)$ is deducted from the mean reversion $q(t)$. This is the only but very important difference with the risk-neutral form of the CIR model. By deducting l you take into account the term premium that the investors expect because of the risk of interest-rate changes. Similar to the previous case, you can calibrate $k(t)$ and $s(t)$ by matching cap or option prices and you can calibrate r_0 and $q(t)$ by matching to bond prices. Unfortunately, $l(r,t)$ cannot be reliably estimated.

This is the classical form of the CIR Model, the risk-neutral and equilibrium model:

$$dr = k \cdot (q - r) \cdot dt + s\sqrt{r}\,dz \;, \quad r_0 = r(0).$$

In this case again look at the difference. If you compare this form of CIR with its arbitrage-free version, you can see that all the coefficients are constants and should be historically estimated. This is the major difference between an equilibrium model and an arbitrage-free one. In the case of an arbitrage-free model you need to match prices for every moment of time t. When dealing with an equilibrium model you take the whole body of available historical data and find such values for the coefficients k, q, and s that feature the best match between scenario statistics and historical statistics.

The last CIR face is the realistically parameterized equilibrium model:

$$dr = k \cdot [q - l(r) - r] \cdot dt + s \cdot \sqrt{r} \cdot dz \;, \quad r_0 = r(0).$$

Again, you have this $-l$ for the term premium and again all the coefficients are constants that are historically estimated.

In conclusion, I would like to provide you with some references that you might find useful:

1. P. Fitton & J. McNatt, "The Four Faces of Interest Rate Model. *Advances in Fixed Income Valuation: Modeling and Risk Management*," F. Fabozzi, ed, Frank J. Fabozzi Assoc., 1997
2. J. Hull, *Options, Futures, and Other Derivatives*, 3rd edition, Prentice Hall, 1997.
3. Lin Chen, *Interest Rate Dynamics, Derivatives Pricing, and Risk Management*, Springer, 1996
4. D. Becker, "Stylized Historical Facts Regarding Treasury Interest Rates from 1955 to 1994." Technical Report. Lincoln National, Fort Wayne, Indiana, 1995

The first paper in the list was the basis for this presentation. I think that it's very well-written and it provides anybody with a very good understanding of what's going on. The second is a classic text. I believe that everybody in this room has one of the editions of this book. If you would like to go more advanced, I highly recommend the third book, but I should warn you that it is not easy reading. However, what is great about this book is that it shows how you can develop a general approach for pricing different types of assets without modeling each of them separately. This general approach is rather complex; it involves stochastic differential equations and partial differential equations, but it's the technique that Wall Street is using right now.

And the last on the list is Dave Becker's paper about his famous stylized facts regarding Treasury interest rates. Sarah is going to provide some insight to this paper. This is an excellent and very practical resource containing valuable observations about the interest rates that were made over probably 30 or 40 or 50 years.

Dr. Sarah L.M. Christiansen: I'm going to be talking about equilibrium realistic generators and why I consider them different. I'm also going to leave you with questions about what you might want in a model depending on the purpose, and especially if you're interested in equilibrium realistic generators.

Now, as I look at it, your arbitrage-free are the top two faces; your equilibrium models are the bottom two faces. The left two faces are risk-neutral; the right two faces are realistic. The guy in the upper left-hand quadrant who is saying " Do it my way" is the finance professor, the bond trader, or the Wall Street guy. The guy in the lower left with all the money is our asset/liability manager, who knows there are times when an equilibrium generator is more appropriate. Up on the upper right we have Madame Olga's Stochastic Scenarios. According to Fitton & McNatt, we need Madame Olga and her fortune telling because you cannot separate the term premium from the credit risk and from just plain noise that's running around. And the guy on the left with the phone and a clean desk is an actuary.

Anyway, we're going to go into more historical background of why actuaries became interested in interest-rate models to start with. Back in 1981 in August, let me read you the yield curve. I don't start with prices; I start with the yield curve. The 3-month rate was 16.76%, the half-year rate was 17.62%, the 1-year rate was 17.17%, the 2-year rate was 16.66%, the 3-year rate was 16.3%, the 4-year rate was 16.08%, the 5-year rate was 15.82%, the 7-year rate was 15.45%, the 10-year rate was 14.95%, the 20-year rate was 13.89%, and the 30-year rate was 13.58%. These are Treasury rates. That is a very high curve and it was inverted by anybody's definition of an inverted curve. Instead of normally being upward sloping, i.e., the longer term has more of a risk premium, it was saying to invest longer you're going to get less. People were selling guaranteed interest contracts to beat the band. And they were guaranteeing rather high rates.

New York has the top reputation as an insurance regulatory state. They were asking some very eminently reasonable questions and these questions revolved around the statutory reserves for the GIC rates. If you have statement reserves in excess of tax reserves and statutory reserves, the rest isn't deductible. So no company wants to hold higher reserves than they have to.

In 1980 they passed an amendment to the Standard Valuation Law (SVL) which said that your interest rate for reserving purposes on an annuity was 3%, plus some weighted factor times the excess of the reference rate minus 3%. The reference rate for annuities is the 12-month rolling average or previous year's average ending on June 30th of the Moody's corporate bond index. That is a long rate.

Now, depending on your plan type, your guaranteed period, or your reserving method, you have different weights. If you have a plan type A that permits no withdrawals, it would be like a bank CD and you couldn't take money out except with penalties. You had to guarantee that rate for 5 years or less, your reserving rate was 12.75%. If you guaranteed it for between 5 and 10 years it was 12.25%; between 10 and 20 years it was down to 11.25% and for over 20 years it was 9.25%. New York asked this very reasonable question: "Can we really assume that 20 years down the road we're going to be able to average 9.25%?"

By the way, who are the regulators? Well, it depends on your point of view and we've all sat, I think, in both seats. If you are a consumer, and I think everyone in this room is, the regulators are protection—they're safety. They are here for us. When we buy insurance or an annuity, what are we buying? We are buying nothing more than a promise; regulators are here to help you make sure that you have a pretty good chance that your promise is going to be met.

The company's point of view is somewhat different. Regulators can be useful. It's always nice to have a good rating and be on good terms with them, but they generally require a lot of red tape and it costs money and adds to expenses to meet the regulator's demand.

New York Regulation 126 was enacted because of these questions that New York was asking in 1982. It said that you'd better do some cash-flow testing, or at that point in time you could use a lower interest rate—in other words have higher nondeductible reserves. Later this was expanded to require cash-flow testing and the law expanded to all annuity products and single-premium whole life, which was at that point in time at least looked at as an investment rather than a protection product.

The New York Insurance Department and the industry got together and they came up with the New York Seven interest-rate scenarios. They are parallel shifts of the current curve with some qualifications. Basically, the first one is that the interest rates you see today are never going to change in the future. The second one is that they're going to go up by 0.5% each year for the next ten years and then remain at that high level. Then you have one scenario that goes up 1% for five years and comes back down 1%. And then you have the one that pops up 3% and then you have the ones where every up is a down, so you go down 0.5% for ten years or then down and then up.

By requiring companies to do cash-flow testing, from year to year, we get some comparability on how our portfolios were behaving. Also they would be able to compare from company to company. There were some minimums—originally it was no lower than 4%; in 1992 it was dropped to no lower than 2.5% on the five-year Treasury and what you did with the rest of it. In the meantime, Larry Gorski, who's another regulator from the state of Illinois and a very good regulator, came up and said, "Wait a minute." You might have an inverted curve and they tend to hang around for about 18 months or so. Also, if the long rate isn't sufficiently larger than the short rate (I believe it's a ten-year rate and the three-month rate), then you have a two-year window in which to get this curve so that it's upward sloping rather than downward sloping.

Again, if you have a company licensed in 49 states, we know it's not licensed in New York pretty much. Anyway, we currently do have an inverted curve and I believe as of this September 30, we are normalizing.

Along with this asset-adequacy analysis, we have New York Regulation 126. A lot of this has been put into Section 8 of the new SVL and risk-based capital (RBC) and United Valuation Systems (UVS) requirements. Every rating agency wants to see what you're doing. Now, New York's Regulation 126 is not limited to those 7 scenarios. It just says you must include those seven scenarios, but you have to do enough so that you are comfortable with the results.

What kind of scenarios do we want to look at? Well, we need some deterministic scenarios. We know that already. The law says we have to include the New York Seven. We might end up with specific tail scenarios for RBC. The interest-rate generator is a very small piece of a much larger, much more complex model. Therefore, some of the software that is required to run these large models is very slow, so you don't want to run hundreds and thousands of scenarios through these models. You are going to be running them through very large portfolios, not just of assets, but of liabilities. So you're going to try to keep the number of scenarios down. We might have some other scenarios. In addition to stochastic scenarios and the New York Seven, we do what we call a static scenario (and shift it up and down). If you took an arbitrage-free, risk-neutral type of generator and you looked at the average because things tended to average out, given the current curve, what would they average to? You can get a single scenario out of that as to how the curve might average over the future. Also, I want to remind you that a scenario is a number of years of complete yield curves, depending on however many you have on your horizon. We generally run 30 years. New York also said in this Regulation 126 that you're running a closed block of business and you're running it out until roughly 90% of the liabilities are gone. Now this can be 20 or 30 years into the future, especially if you have something such as structured settlements, or you might have purchased a retirement liability for a company that terminated a pension plan, and they have employees age 25 who are not going to start collecting their annuity benefits until they're 65. Well, that's 40 years before the annuity starts, so you have all sorts of liabilities, you have annuities in payment, you have all sorts of things that are going on that can be very long, and other things that may be quite short.

We're not going to try to establish a fair value of an asset or a liability. We're going to try to determine whether our reserves are adequate. Solvency is perhaps too strong a statement, but basically we're going to do some testing on consumer protection and we're looking at a very, very long time horizon. And unlike a securities firm where the litigation risk disappears, very shortly, they're not likely to get sued for something that happened 20 years ago. Those of us, who remember the vanishing premium scheme, know that we do get sued 20 years down the pike for something that happened. And the fact that you put some sort of statement in your illustration that's comparable to "past performance is no guarantee of future results" in a prospectus for a mutual fund is not sufficient protection. That doesn't cut a whole lot of weight, so it doesn't prevent us from class-action lawsuits.

So what kind of model do we need? Vladimir and John started in the arbitrage-free corner. Since modeling is a method of successive approximations, you take your model and then you extend it a little bit. They came down to the bottom corner and they had models with stochastic differential equations. Well, I have news for you: Nobody who started out in the bottom corner ever thought about stochastic differential equations. We didn't go talk to the physicist. The early models, those of the late 1980s, Jetton and Strommen, Gurski, Mereu, myself, all of us started out with models that were not arbitrage-free with respect to the current curve, and they were all basically using discrete time steps. Time steps tend to be at least a month, often three months. Ask yourself, is continuous time important if you're looking at a time step of three months?

We also have a small number of factors. Jetton and Strommen used a single factor model based on 20-year rates. A lot of the original models did use short rates. Gurski used three sets of rates; Mereu used inflation one- and ten-year rates. The Markov Chain Process used a 20-year rate and a shape code. By the way, that's been improved since that paper. All of these involve nonparallel shifts of the curve and produce stochastic scenarios.

There are differences in the way that all the intermediate points on the curve were determined, whether it was from a deterministic formula or not. My shape codes were each associated with factors and you multiplied that by the long rate. You came up with a tentative rate. And then there were some additional constraints, and those constraints could be the difference between the long rate and the short rate. They could be how much your rates could change from year to year or quarter to quarter—minimums, maximums, other spreads and other such things. You want to keep it simple. You want to minimize the complexity and you want to use as few independent variables as necessary to specify the model.

David Becker and Mark Tenney together, apparently, are responsible for this collection of stylized facts about U.S. interest rates. Dave is the one who wrote the paper, and I was told at a seminar in New York recently on interest-rate models that Mark Tenney was at least as involved and it was really his idea, so we need to give Mark some credit.

U.S. interest rates are not negative. You don't want this in a model and there are some of these stochastic differential equation models out there which do produce negative rates. Interest rates don't go to zero and stay there nor do they head up toward infinity and go high and stay high, but rather they remain within a reasonable range.

Also, your volatility tends to decrease with increasing maturity, which is why, while most of the models depend on the short rate, my model keys off the long rate. There was a very good paper recently in the July 1999 issue, Vol. 3 No. 3 of the *North American Actuarial Journal* titled, "Term Structure Models: A Perspective from the Long Rate." The author also found out that inversions tend to occur about 15–16% of the time and there are varied definitions of inversion. I like to say that a yield curve is inverted, when for the most part the curve is monotonically decreasing or at least not increasing. A lot of the definitions of an inversion only look at two rates, say the ten-year and the three-month, and if the three-month is above the ten-year, you have an inverted curve. They have asked what if the 3-year were above 105% of the 10-year? That cut the inversion down to about 11%, and those 11% of the time tended to be times of economic stress, political crises, and other such things.

What's important is that you want to look at all of these things when you're looking at a model of asset-adequacy analysis. Fitton's paper called this a wind tunnel for actuaries, and that's really what it is. We're doing some very heavy stress testing. We're going to produce a complete yield curve for each discrete time step of our time horizon of 30 years or more. We're not just going to produce Treasury rates. We're going to produce portfolio rates and Treasury rates. We're going to do bond rates; those of us who have commercial mortgage will have commercial mortgage rates. We might have rates for residential mortgages and ideally our spreads will vary by scenario, by maturity, and by year in the future.

It is true that the spreads do not stay constant. There are lots of times, where if you have a very wide spread, it may narrow in the future. I don't like some of these software models out there that say, now I can vary my rate by maturity and by credit quality, but I have to keep it the same for all years in the future. What if my bond rate spread was 200 basis points? Does that make sense that it's going to stay that high? Does it make sense that it would stay at 30 if it happened to be 30 for 1 point on the yield curve? No, they're going to vary. We want to be able to evaluate assets and liabilities on a common scenario. We use average portfolio rates; we include spreads and our spreads are implicit. We generate all the scenarios for each of the different kinds of assets, blend them, and let the chips fall where they may. Although we do some checking on bond spreads to start with. They may be public issues or private issues and while we're blending, we blend based on an accumulated value of one and turn that back into the interest rate for that point on the yield curve.

We'd like to have a scenario generator that's robust for which we provide a variety of interest-rate curve shapes, including upward sloping and inverted, as well as humps and valleys. Humps are quite common, but valleys are exceeding rare, although it was not very long ago, one time in this past year, when it seemed to me that we had a valley. Level and perhaps oscillating. This curve looks like it's almost level, but it's sort of wiggling around and it might look a little bit like a sine curve or a cosine curve or something like that.

We'd like to be able to get results that are comparable to historical results. We don't want to have to recalibrate frequently. In fact, the less frequently we recalibrate, the happier we are. Because we want to be able to compare this year's results to last year's results, which are the previous year's results and, if you're re-calibrating your generator every three months, then there's nothing in common between this year and a year ago.

Let's take a look at science for just a brief minute. Science knows that you don't use more significant digits than you can measure, so let's ask ourselves—what are we measuring? We're looking at the impact of the change on interest rates on asset behavior and on policyholder behavior and how they relate. Now, do we want to have an absolutely stupendous perfect model for interest rates and almost no information on the way that the liabilities react? We may have pretty decent information on asset behavior; sometimes we do and sometimes we don't. But how many of you have done lapse studies tying your policyholder behavior to your interest rates? Very few. So you need to look at that. I mean death is independent of interest rates. We can say a person isn't going to die depending on interest rates, but I may well lapse depending on interest rates. Or he or she may put more money into the product depending on interest rates. So we want to make sure that we are not looking at things and saying, we have this wonderful, super-duper model that we got from XYZ Corporation and it's absolutely the perfect interest-rate model. It does everything everybody wants, but we have no idea how some securitized and collateralized debt obligation behaves with respect to interest rates. We need to do some policyholder studies, and we need to do some experience studies work along with it. If all I convince you to do is to start collecting the data so that ten years down the road they have some data to look at, that's a great start.

Again, we're not going to try to predict rates. We'd like to have a reasonable amount of time to run the entire model. We don't want to run more than 100 scenarios. That seems to be a cut-off that most companies want to work with and we want to provide ourselves some useful information and we're much, much more interested in the tail of the scenarios and where we are having problems than we are on the means.

If you're pricing an asset for sale or you're pricing a block of business for sale, you'd better be using an arbitrage-free scenario so that you can be with the market. If you're doing stress testing or asset-adequacy testing, you're not interested in the mean; you're interested in the tail. When are you having problems and why? We're looking at a distribution of surplus. Even though we'll start with a statutory value of assets equal to the statutory value of liabilities, we're going to be comparing the "market value" of assets to the "market value" of liabilities. What we really are comparing is the present value of the future asset stream to the present value of the future liability stream under each of these different scenarios for each of these different times. There may not be a market for everything, and only at times zero would have you have any sort of real market value.

The theory, when Regulation 126 was originally established, was that the New York Seven would give us our worst-case scenarios. That's been close to true, but not entirely true. We have generally found our worst case happening on our stochastic scenarios. However, in the worst ten or so generally one of the New York Seven has been in there. We'd like to be able to get the reliability of running 1,000 scenarios, when we're only running 50. We've written some papers on representative interest-rate scenarios, and let me tell you, I know it sounds like going to the moon, but we did get to the moon.

The relationships between products and interest rates vary. Some products have problems with rising rates, others with falling rates, and many of them with inverted rates. Some have a level or a humped curve, especially a hump because the back half of the hump looks like an inversion. We'd like to have a relationship between equities and interest rates. We really need to do some lapse interest-rate studies. The amount of time it takes to run a single scenario through is critical.

So let me leave you with a batch of some unusual questions. These are the questions that I'm asking myself. Does a continuous time model make sense if you have a three-month time step? Does a complex model that requires quarterly recalibration make sense if you're doing adequacy analysis? So those are the questions that I had and I hope that this has been of interest to you.

Ms. Albert: Well, we've had quite a bit of information here. Did anybody have any questions that they wanted to ask of our panel?

Mr. Jeffrey S. Roth: Sarah, you mentioned about the best way as far as with spreads, you should not assume a constant spread; that it should vary. I was wondering how you go about doing that?

Dr. Christiansen: Well that's a really good question. We do it by taking our model and running the corporate bond curve for, say, our average quality through that same model with the same "pseudo random numbers," we'll call them for purposes of clarification, and see what we get for bond rate scenarios. Run through again with the same random numbers and the Treasury curve to obtain Treasury scenarios. Now, if you subtract off the Treasuries, you get spreads once you convert everything back to yields to maturity curves. Then you really should check that these spreads are not out of bounds. There's probably some historical data on what these spreads are, and that's another thing. We're starting to track what our spreads are in-house and to develop a database on spreads, so we have some idea of what's a reasonable spread and what isn't. We have had to do some adjusting and since our model, in-house, works primarily off a blending curve and does not work off Treasuries, I tend to adjust our Treasury model so that spreads stay within reason. There is another area where some of our people use PTS where they do add a spread to Treasuries so I think there are a couple of different ways. We've gone about it from an implicit way. I suppose also if you got some good statistical data on what the spreads were, are your spreads high when your Treasuries are high or are your spreads low because the Treasuries are high? You could do it either way by getting some statistical data and modeling it directly, or we go about it backwards and back them out.

Mr. John M. Bragg: I'm certainly glad that Sarah brought out the vanishing premium situation because it certainly proves that we get into a lot of trouble on the nonguaranteed elements as well as the guaranteed elements. And, of course, as Sarah would agree, the entire universal life scene in the early 1980s was part of that also, with tremendously high interest rates being estimated for long periods in the future that got us into an awful lot of trouble and I guess we're still in it. That's just a comment. The question, really that I have, maybe to John and Sarah both is, what accounts for this strange yield curve situation this year in 2000 and, in particular, do the panelists think that the Federal Reserve moves have anything to do with it?

Dr. Christiansen: Well, I'll try to comment on that and then I'm going to pass the buck over to John to correct my comments. First of all, the Federal Reserve can affect the short end of the curve. Secondly, we have the Secretary of the Treasury selling long bonds because we have pulled ourselves out of debt. While they have a chance they are refinancing into shorter terms, so they're buying the 30-year Treasuries, which is increasing demand and decreasing supply. This is lowering the 30-year rate and that also goes with the fact that if you have been following the news, that our bellwether rate is now the ten-year Treasury as opposed to the 30-year. Beyond that, you can probably talk about what the market expects and, of course, there's the expectations theory of the term structure which says that basically they think rates are going to go down. And there's also a liquidity premium theory and preferred habitat theory, and with that I'm going to turn it over to John who knows more.

Dr. Marcsik: Actually, I don't have too much to add to that. I think that was a very good response. Only to add that just because a yield curve is inverted does not mean the market is not implicitly charging a premium for longer bonds. It's completely consistent to think that if the Federal Reserve is raising short-term rates and it's a cyclical problem that will go away in a few years, that the market would believe that short-term interest rates would fall again once the cyclical problems go away. But their expectations might be lower than they are today. You can still add a positive number to that and end up with an inverted yield curve.

Ms. Albert: You're adding smaller positive numbers than you are to the short term.

Mr. Larry M. Gorski: I think the history lesson on the development of the valuation law was very good and maybe to point out that in the not-too-distant future the required interest-rate scenarios may be dropped. The requirement for New York Seven may actually be a thing of the past. At session 88PD "NAIC Actuarial Opinion" that I'm moderating, there's going to be a discussion about the proposed revisions to the Actuarial Opinion Memorandum Model Regulation. And that's one of the things that is on the plate, to drop that required seven. And because of that direction, I think sessions like this are very important to spread the word about all the elements of scenario generation (interest-rate generation) that should be considered. Because, if in fact, what I just said does take place, the so-called safe harbor of the New York Seven won't be there any more, and actuaries are going to have to take responsibility for the decisions that they make with respect to interest-rate generators and cash-flow modeling in general. I also think Vladimir's comments concerning the risk-neutral parameterization were very good. I've always had a tough time trying to understand that and in two sentences you got the point across pretty well.

Ms. Albert: Larry, in terms of going away from the New York Seven, is that because when companies use those seven scenarios they generally flunk one or two of the tests?

Mr. Gorski: I think the general reason is that most regulators feel that the New York Seven isn't sufficient to really test the strength of the company's asset/liability management process. In my regulatory role I'm oftentimes involved in the asset-related issues and I see how structured some of the assets are and how they react in different interest-rate environments. The New York Seven scenarios with their parallel yield curve shifts don't really get to those issues. I think that we all feel that more testing needs to be done beyond the New York Seven. Yet people look at the New York Seven as a safe harbor; if we do that, we're done with our testing. So the feeling is that if we remove the safe harbor and put the authority and responsibility in the actuary's hands, they'll do a responsible job. And so it's not because companies are failing one or two tests; it's that we do want to see actuaries doing a complete job in this area and not just looking at the New York Seven and opining on that basis.

Mr. John B. Gould: I just wanted to ask the panel generally, with more and more of the development of new products that are dependent upon such things as S&P indexes or perhaps in the case of European mid-term notes and various types of European securities, there has to be obviously a consistency between the interest generator you have and, of course, the generators that you're going to use for all the indexes. This is going to involve correlations, cross-correlations, and probably cross-cross correlations. At what point do you feel that the system breaks down, or do you think it should be extended that far?

Dr. Christiansen: Well, John, I'm going to take a stab at this. In Dave Becker and Mark Tenney's research they came up with the fact that there was not a significant correlation between the equity markets and the interest-rate markets. On the other hand I'm going to put on a different hat and say I'm chair of the Finance Research Committee and we are looking for a member on a project oversight group (POG) to help look at this particular problem. We have a call for papers that we're ready to put out when we get our POG formed and get this written.

Dr. Ladyzhets: I just want to mention there is a 5-factor model, which is part of PTS, which does have S&P 500 together with interest rates. And the basic idea is that this model does not use correlation but it uses a connection between volatility of S&P 500 and interest rates. And I believe that this is a connection based on Dave Becker and Mark Tenney's stylized facts. But about correlations—it's hard to say when you're going to break it. Definitely the bigger correlation matrices you have, the more headaches you have. For example, sometimes it's very difficult just to check that the correlation matrix is positive definite. So the idea is that at some moment you should choose which factors are the most important and build your model from them. My feeling is that if you go above ten factors, you will be lost. That is just my feeling from working with this model.

Dr. Marcsik: I think there's always a danger, especially when you're historically calibrating these models, to do what statisticians might call data-mining. If you take enough statistics of history some of them will be inaccurate, so you don't want to use the finest detail of every statistic. Also, statistics is applied in this type of modeling—it's not ideal. If you're measuring the way elementary particles work in particle physics, the distributions are very stable. It doesn't matter if you do it on Wednesday or Thursday; you always get the same distributions because you've put yourself in a laboratory condition where you can make precise measurements under exactly the same conditions—at least exactly the same for anything that matters in your measurement. All of these markets are measured; prices are measured in a constantly changing environment. There are always different politicians; there are always different things going on in the economy: political, global, whatever. So I would say you should always be cautious when you're using positive or negative correlations and I would definitely advise you to stress-test them. I would be very leery of anything that depended a lot on an exact correlation, so if you were trying to price some guarantee on a basket of assets, that historical correlation could be a statistical accident because of some reason that no longer exists and won't be there in the future.

Ms. Regina Lisa Lefkowitz: Well, since Dr. Ladyzhets brought up the PTS 5-factor model, I feel free to ask him this question. Were you expecting mean reversion to come out of your model?

Dr. Ladyzhets: Yes, basically we were.

Ms. Lefkowitz: So obviously we just didn't run enough scenarios to get it.

Dr. Ladyzhets: No, the number of scenarios should not affect your ability to see conversion to the mean. But what is really important is how long your scenarios are in terms of the speed of mean reversion. There is so-called time for mean reversion, which equals to one over the speed. If the length of scenarios exceeds this time, you will see the effect of mean reversion. However, if the scenarios are too short, you just will not have enough time to observe it. For example, if your speed is about 0.001, you need a significant time in order to feel the effect of mean reversion. What is going on is the following: you have mean reversion, which is returning your interest rates to some level, and you have the stochastic part, which is trying to change the interest rates stochastically. They fight one another. That is why you need some time in order to have the mean reversion overcome the effect of stochastic deviations.