RECORD, Volume 26, No. 2*

San Diego Spring Meeting
June 22-23, 2000

Session 38PD
Financial Modeling Integration

Track: Investment

Moderator: JOSEPHINE ELIZABETH MARKS
Panelists: CRAIG MERRILL†
RUSSELL A. OSBORN
Recorder: JOSEPHINE ELIZABETH MARKS

Summary: While actuaries have developed long-term modeling techniques using cash-flow analysis and stochastic methods for insurance product pricing and valuation, financial analysts have focused on discrete mathematical solutions and option-pricing techniques for derivative pricing and valuation. With the recognition of insurance products as the "long-term derivatives of the financial services industry," there is an expectation that these techniques may be consolidated into new tools for financial analysis and modeling.

The panel discusses:
- *Modeling tools used by the actuarial profession*
- *Pricing techniques used by Wall Street financial analysts*
- *Advantages and disadvantages of each approach*
- *Market implications of each approach*
- *Potential for harmonization of the two techniques*

Ms. Josephine Elizabeth Marks: Our first speaker is Craig Merrill. Craig is a professor of finance, and a fellow of insurance and risk management at Brigham Young University. He holds a Ph.D. from Wharton, and has published various papers, including a monograph, on financial valuation models.

Mr. Craig Merrill: This topic is one I've been thinking and writing about for quite awhile now and there's really an awful lot to say about it. Rather than try to cover everything, I've picked three general areas to make some comments on. My goal is to build intuition. The technical detail is far beyond what we can do with this amount of time, but there are a number of references available for the technical details. If you are interested in having a copy of this presentation, it is available on my home page at http//arbitrage.byu.edu.

The first thing I want to talk about is the financial pricing of securities. The overall session topic has to do with integrating financial modeling techniques into insurance

*Copyright © 2000, Society of Actuaries

†Dr. Merrill, not a member of the sponsoring organizations, is a Professor of Finance and Grant Taggert Fellow of Insurance at Brigham Young University in Provo, UT.

company problems. We're focusing mostly on the liability side, but I want to talk a little about financial pricing in general to point out a couple of mistakes that get made, places where people get easily confused or misunderstand what's going on.

I've constructed a simple little example that highlights the way that we deal with risk when we're pricing in the financial paradigm. The first observation is that investors are risk adverse. Because investors are risk adverse there's a risk return tradeoff that has to be acknowledged in pricing.

The level of risk aversion differs across investors. That allows markets to function. If we all had exactly the same level of risk aversion, we'd get to a market-clearing price, and there would be no more trading. Because we disagree about estimates of risk and expected return, and we have different levels of risk aversion, there can be an active market with trading. The prices have to reflect the fact that investors have varying degrees of risk aversion, and the market-clearing price represents something about the aggregate, or average, level of risk aversion in the marketplace.

There are three general approaches that are used in financial pricing to deal with risk in order to come up with a price that reflects what typical investors would be willing to pay. Here's the simple example. We've got a security that's got one period left till maturity. At maturity, it's either going to pay $110 or $90.

One approach would be to add a risk premium for discounting, to reflect the uncertainty about the ultimate pay off. We assume 50% probability for each cash flow then we discount at a risk-free rate of 5% plus a 2% risk premium. We're discounting at 7%, and we come up with a price of $93.46.

Now the risk premium is what captures investors' attitudes about the uncertainty in the future cash flows. If investors were not risk-averse, that wouldn't be there. All we'd have is just the time value of money portion, and we'd have a price closer to $95.

Another possibility is to change the probabilities. Now, this goes under a variety of names. You might have heard of risk-neutral probabilities, Equivalent Martingale Measure, risk-neutral pricing, and arbitrage-free pricing. The notion here is just the same as in the first approach. We've just chosen a different number to change. Instead of changing the discount rate to reflect risk, we changed the probabilities.

There are many mathematical theorems that tell us why and how, and under what conditions, we can do this, but the point is this: done properly, what we're doing is a change of measure. We're changing from one probability measure to another, both of which yield accurate pricing. We adjust the probabilities in such a way that when we discount at the risk-free rate we come up with the right price, the same price as in the last approach. If we don't come up with the same price there's something wrong with our pricing model, our technique for using this method.

Now, typically, the arbitrage-free or the risk-neutral pricing methodology gets applied with derivative securities. If you think about an option on a stock, the option has value only through its relationship to the underlying stock. If we have an active trading market in the stock, and we've got an active borrowing and lending market, then what we can do is construct a portfolio of the stock and borrowing or lending that exactly replicates the option payoffs over the next time increment.

If we construct a portfolio of two actively traded instruments, you can think of it as money market and the stock, that exactly replicates the option, then the price of the two has to be the same. We have a value for the option. Out of that hedging, portfolio construction falls. Mathematically, these probabilities are called a risk-neutral measure or risk-neutral probabilities. There's nothing intuitive about them, per se, but they work. The reason they work is that essentially we're using an adjustment to the probabilities to account for risk aversion in investors.

There's nothing that says that changing the discount is theoretically preferable to changing the probabilities as long as we do either one right. There are situations where one approach is easier to use than the other. For example, it's easier to price options using risk-neutral probabilities, but it's easier to do corporate finance using the discount rate because we don't have the ability to construct the hedge portfolio, and back out the risk-neutral probabilities.

Well, we've got three things in our price calculation: cash flows, probabilities, and discount rates. We've talked about changing the discount rate, and we've talked about changing the probabilities, that leaves one other. We can change the cash flows. Instead of having $110 or $90, we say what one single cash flow, received for certain, is just as valuable to the risk-averse investor as the risky cash flow opportunity. That's called a certainty equivalent cash flow.

In this case, $98.13 in this market, would be equally desirable to risk-averse investors as facing a lottery between $110 and $90 with equal probability. Now this $98.13 is less than the expected outcome of $100 because of risk aversion. We're willing to give up something in order to get rid of uncertainty because we're risk averse.

We've got these three methods for dealing with risk. We can adjust the discount rate, we can adjust the probabilities, and we can adjust the cash flows. Again, this is very important, done properly, they all give the same answer because in a well-functioning market there can only be one price, or else we have an arbitrage opportunity. If our model is yielding one price with one tool and another price with another tool, we're doing something wrong in our modeling.

Let me jump to interest-rate models now and then we'll have some specific examples of where the first part of this comes into play. Interest rate models are constructed in the following manner. We make an assumption about what the fundamental driving sources of uncertainty are in the marketplace. The simplest models assume that there's one source of uncertainty. For example, the shortest maturity spot rate, which we sometimes call the short rate, would be modeled as a stochastic process. We know its current value; we know how it evolves over time.

There are more complex models that would have the volatility of the spot-rate process itself be stochastic or the level that the mean reverts to (if we have mean reversion) could also be stochastic or both. Or we could have multiple spot rates being stochastic, or forward rates being stochastic, which is the Heath-Jarrow-Morton Model.

We make an assumption about how many sources of uncertainty are needed to capture the stochastic environment of the interest-rate market that we're in, and choose some sort of stochastic process for each of those sources of uncertainty.

We then have to make an assumption about equilibrium in the marketplace. That equilibrium could be a term-structure theory, like an expectations hypothesis of some sort, or it could be an absence of arbitrage type of equilibrium assumption.

Once we've assumed the number and type of stochastic processes that capture uncertainty in the marketplace, and the nature of equilibrium, the rest of it is mathematics. Now that's not to say it's easy, but it's mechanical because everything's determined once you've made those first two assumptions. You plug the mathematics, and out comes your interest-rate model so that bond prices are a function of maturity, the yields on zero-coupon bonds are your spot rates, and you can infer forward rates from the spot rates, and so on.

Now in terms of having them fall out mathematically, there are two general approaches that are used. One is closed-form solution, which for one- and two-factor models have been fairly forthcoming, for three-factor models have been quite difficult, and for some of the more complex derivative securities have been very difficult. As an alternative, if there's not a good closed-form solution where we can just write down a formula for the interest rate or the security price, we use numerical methods.

If you want to generate a term structure, for example, simply doing bond prices, a closed-form solution gives you the whole term structure instantaneously. But if you've got to do numerical solutions and you want good solutions, you're going to need about 10,000 simulations for each price, and it's going to take you on a reasonably fast computer 20-25 minutes to generate a term structure. Then if you want to move that through time, using the numerical solutions can be very costly, but it's doable. We can do some really fairly amazing things with the modeling technology that's out there.

Now there are generally two types of models that we look at in the financial modeling area. One type is what we call an equilibrium model, which gives an equilibrium condition related to the risk aversion of investors in the marketplace. It tells us about the risk-return relationship of all interest-rate contingent securities at a given point in time in the marketplace. We make an assumption about risk-return relationship, and what the market price of risk is, and that allows us to then derive bond prices, interest rates, and derivative security prices.

In the other type of model, we assume the absence of arbitrage, and where markets are complete, (i.e., we can replicate any security with a set of fundamental securities) we're able to derive prices (especially of derivative securities) by assuming the absence of arbitrage.

Arbitrage-free gets used in two ways. Sometimes I think this can be a little bit confusing. One use of this term, arbitrage-free, is in defining equilibrium. We want to have a model that in equilibrium does not admit arbitrage opportunities between the various securities in the marketplace. That's an economic concept.

The other use of the term arbitrage-free is in trading and hedging activities in the models that we use on Wall Street for trading decisions. In that case, what we want is a model that exactly replicates all of the security prices in our data set that we're using to calibrate the model. If we've got an entire term structure of bonds, we want every one of those bonds to be exactly repriced by our model so that the model does not admit any apparent arbitrage opportunities at that point in time.

If you consider discrete interest rates that are observable in the marketplace, an arbitrage-free model would fit a term structure that goes exactly through each of those data points, so that at that point in time the model does not appear to admit any arbitrage opportunities.

An equilibrium model, on the other hand, or an arbitrage-free model in the economic sense of arbitrage-free that has a limited number of parameters, is going to be a fit to scattered data, more like a statistical fit. That's not going to be ideal for trading and hedging activities because it seems to say that there's an arbitrage opportunity between the model price and the market prices. But that can't be right because we take the market prices as correct.

We've got to make a decision about which type of model to use: the one that fits very tightly or the one that captures the general shape relationship. Well, the relative benefits of the two models can be seen as you go forward a little bit in time and your market data changes. Now, which has the largest errors relative to the new market data? Statistically, it's likely to be the arbitrage-free model. It no longer fits the data perfectly, but our equilibrium model has a nice relationship to the new interest-rate environment.

Each of these modeling approaches has strengths and weaknesses. Equilibrium models tend to be really good for interest-rate scenario generation. We want to capture a general interest-rate economic environment and project it out into the future. In that case, we want the general shape. We don't want the specific jagged peaks and valleys from today's data. When projecting out into the future, the errors relative to future interest rates tend to be smaller with this type of model.

It's good for risk management. Again, it captures the general interest-rate environment, and is more likely to allow us to make good decisions about interest-rate risks that we're taking in our investment portfolio or in the liabilities that we're issuing. It also tends to do a good job when we're making long-term investment decisions. For instance, when we want to find specific securities that are under or overvalued relative to their intrinsic economic value, as opposed to today's idiosyncrasies in the interest-rate market.

Now, on the other hand, the arbitrage-free models tend to be very good for hedging decisions. If you're issuing interest-rate options and you need to make hedging decisions in your trading activities, they tend to be better models to capture today's specific environment. They're good for short-term trading decisions. They're also good for inter-market arbitrage trading. If we're looking at one market and then asking if another market is priced correctly relative to the first, the arbitrage-free models tend to be more appropriate.

Again, both the simpler closed-form solutions and the more complex, robust models have to be solved numerically. When we're doing trading type of activities, the numerical solutions become a constraint in finding trading opportunities. The time that it takes to process the model can make the arbitrage opportunity obsolete by the time we discover it. You see the investment banks invest in parallel processing super-computers. Whereas in an insurance type of environment, more often we tend to be interested in the first set of objectives and the numerical solutions may be less of a problem. Though, if we want to have lots of interest-rate scenarios, it does become costly.

The third thing I want to comment on is this issue of valuing insurance liabilities. Now, I've talked a little bit about financial pricing and the approaches to dealing with risk, and I've talked a little bit about interest-rate modeling. Ideally, what we'd like to do is figure out which of the three approaches to financial modeling and which type of interest-rate models combine to best allow us to value insurance company liabilities.

But there are some real challenges that we need to deal with before we even do the mechanical part of it. I believe that the implementation of interest-rate models and choosing which pricing approach to use, is a fairly readily solvable problem. I'm not saying it's easy, and as you'll hear from Russ, there are some significant challenges in implementing some of these tools in the insurance area, but there's not a lot of big theoretical hurdles to that idea.

There are some real conceptual problems, though, in deciding just where we need to apply these financial pricing techniques. I think that one of the big ones is in defining just what insurance liabilities are. If you've seen articles over the last few years in the *North American Actuarial Journal*, there's been quite a discussion of valuing liabilities and valuing various approaches that might be used including option pricing, corporate finance techniques, and traditional actuarial techniques. But the thing that's always confused me as I've tried to follow that line of literature is just what they are calling liabilities.

Some articles are looking at a block of business. A block of business meaning policies that have been issued and the assets that back them, which typically is what changes hands when you look at mergers and acquisitions, or trading of liabilities. It's the entire block that moves, and there is value in figuring out what the market value of that combined entity is, but that's sort of unusual from a corporate finance or an accounting point of view to call the assets and the liabilities that go together, liabilities. Typically we look at the cash flows that are obligations of the company, and we call those things liabilities.

What is it that we're talking about? Is it the block of business? Is it the projected cash flows? Well, here's where the difficulty comes up. If it were just guaranteed cash flows, it would be a pretty easy thing to solve, but insurance companies have non-guaranteed cash flows, which depend on the asset performance as well as the decisions of the investment committee, and whatever else might come in to play: actuarial experience and expenses. Projecting the cash flows can be challenging.

There are also issues that surround the present value of cash flows. Do we project them using risk-adjusted rates? Do we discount them using risk-adjusted rates? Should we project them and discount them using risk-free rates? Is it some combination of the two? There are some real challenges in this area, and some things that we need to try to answer.

To try to gain some insight I want to draw a couple of analogies. The first one is to mortgage-backed securities. When mortgage-backed securities first came into being, prepayment risk was a huge challenge. How do you deal with prepayment risk? People don't prepay in an economical optimal fashion. People make decisions based on being transferred by their company. If I'm transferred, I prepay my mortgage because I'm not going to keep a mortgage on a home I can't live in. People make decisions based on divorces. They make decisions based on a conviction that they're going to get out of debt. They might make decisions based on interest-rate changes, which is what we think of as economically rational. Here you are trying to value mortgage-backed securities, and you've got this economic irrationality in the way people behave.

Well, a number of models have been developed that try to link prepayment behavior to interest rates even though it's a noisy relationship. What we've found over time is that applying the same sort of arbitrage-free pricing models that we use in more complete markets, has been workable. In fact, as the market has evolved, the option-adjusted spreads have shrunk, and we've been able to either very clearly quantify what the risk premium ought to be for risks that we're not capturing in the model or completely remove it, because we're now capturing the prepayment risk more effectively through the models themselves.

We've got the same sort of challenge with insurance. We've got crediting rates, policy loan options, and policy termination options, all those things like the prepayment option on the mortgage, are non-traded. But application of the mortgage-backed security pricing techniques is likely to render these things workable over time. It will take some time for the data in this area to emerge, and for option-adjusted spreads to stabilize or disappear, as the modeling technology matures, but it ought to be workable in a similar manner to the mortgage-backed security market.

The key point here is these things have to be done regardless of which interest-rate model or which equilibrium definition we choose to apply. This is a challenge that needs to be tackled regardless of which financial pricing paradigm we choose to apply to insurance liabilities.

What are the remaining issues? There's a lot of thought and effort that needs to go in to deciding what are the appropriate discount rates for projecting cash flows. The real challenge in this area has to do with the interdependence of liabilities and assets in the insurance company setting. The uniqueness of our measure of liabilities ends up in accounting statements, and ends up being used for management decision-making.

Think about two companies that have both issued exactly the same sort of variable life policies so that the composition of policyholders is very similar between the two companies. Now they're both going to pay crediting rates based on the investment strategy taken. One of them is using AAA bonds to back these policies. The other is using pork belly futures. Quite a difference in risk!

Now, if you project cash flows using the expected return of the investments, the cash flows for the pork belly-backed life insurance policies are going to be larger than the ones backed by the AAA bonds. Higher risk, higher expected return. Then if you discount them back at an appropriate risk-adjusted interest rate, the risk-adjusted present value of the two blocks of business will be at about the same level.

If you then swap the pork bellies into AAA securities, and keep the pork bellies discount rate, then you'd see the pork belly-backed life insurance policies change significantly in value. The assets that back the block of insurance contracts have an impact on how valuable that block of business is. If they start out at the same value, and then you change the assets that back them, they can change in value if your methodology doesn't carefully account for the risk that's involved.

Market participants such as regulators, investors, and mergers and acquisitions analysts, understand market valuation of a block of business. What they need in financial reporting is more transparency of the components that go into market value.

Let's think about corporate bonds as an analogy. You can have two companies issue bonds that have the same promised cash flows, but the bond values can vary depending upon the assets and the business opportunities that are backing the two bonds. You can think about a corporate bond's market value being a risk-free value, minus the put option, the right of the company to default on the bond. The more risk the company's taking, the larger that put option will be, and the lower the market value will be.

The same thing holds true with insurance companies. We can project the cash flows out, and discount them back to get the risk-free value. Then we can look at the put-option value, and that becomes a measure of how much asset risk the insurance company's taking in the way it's doing business. That we could actually try to quantify the asset risk that the company with the pork belly futures is taking relative to the asset risk that the company with the AAA bonds is taking.

The bonds have promised cash flows. Now those cash flows may or may not happen, depending on default, but we know if they happen how big the cash flows would be. This is where we have the real challenge in the insurance area. We don't know just how big the cash flows will be because there are non-guaranteed cash flows. This is the area where we're trying to figure out what ought to be done.

What's the appropriate mix of discount rates and risk adjustment for projecting cash flows, and then for discounting them back? Two things need to be accomplished. One is internal management, and the other is the observation from the outside, whether it's investors, regulators, or mergers and acquisition type of analysts. Those people want to have some sort of transparency into what the company's doing.

My feeling is that some of the applications of financial theory are relatively straight forward, once we work out some conceptual challenges. These challenges are interrelated both with valuation and financial reporting questions. We don't have time to go into detail but I hope to have clarified some of the intuition behind the thought process in financial modeling, some of the issues in application to insurance, and some of the challenges that are still remaining.

Ms. Marks: Our next speaker is Russ Osborn. Russ is a senior actuary at Nationwide with extensive ALM experience. His responsibilities include bringing together product strategy and investment policy, and also leading the integration of software and business applications at Nationwide.

Mr. Russell A. Osborn: The main objectives of my presentation are to build on what Craig has just talked about, and leave you with a good understanding of the sophisticated pricing techniques of both actuaries and quantitative financial analysts on Wall Street. When I say sophisticated I am referring to the actuarial techniques using stochastic models that look at lots of scenarios, although I suspect that some companies are still using a best guess scenario for doing their pricing, and that's just not going to cut it in the new age.

We want to talk about the appropriate uses of the two models, advantages and disadvantages of each, and whether or not they can be integrated in certain situations. To do this there are really a couple of key insights that you need to come away with. The first one is to realize that there are differing goals between investment banking and insurance.

Investment bankers have a very short horizon. They're in the business of using as little capital as they can to turn deals over, to find a buyer and seller, and to collect a fee for matching them up, but otherwise take as little risk as possible. The pricing models that they developed are designed with that in mind, to find a model that buyers and sellers can agree on as a way of communicating prices.

On the other hand, an insurance company is going to look at a risky vehicle or a risky investment from the perspective of actually underwriting this risk, and keeping it on their books for the long haul. We need to understand the differences between the risk-neutral models and realistic models that are used in actuarial appraisals for making decisions. Do not confuse the distinction between risk-neutral and realistic models with the distinction that Craig made between risk-neutral and equilibrium models.

I'm going to use a similar example to Craig's, where we start with a simple asset with a price of $95 in an actively traded market and there are two states of the world. We can impute mathematical probabilities to the particular states of the world to express our price as an expected value.

We can take the first outcome, $110, multiplied by the nominal probability of 25%, and weight the $90 with a 75% probability. If you're given a market price of $95, and you have a particular model of the world that tells you the possible outcomes for that asset, then you can backwards solve for these risk-neutral probabilities that allow you to express prices as expectation.

Now I want to distinguish that from what I'm going to call a realistic model. In this case I would like to assume that your company has some sort of expertise in the underlying risk that determines whether there's a $90 or $110 payoff. I also assume that you have looked at the risk underlying this payoff, and maybe have done historical studies or econometric studies.

For example, in the case of earthquake risk that was used to underwrite catastrophic bonds, you would have a staff of experts in geology and earthquakes. If your expert team came up with probabilities of 50-50 for the two outcomes, then your expected value using these realistic probabilities would be $100. That's because these were not designed for coming up with market values. They were merely designed for giving a decision-maker some information about the distribution of outcomes, and the expected value of the payoff.

What any particular investor would pay really depends on their individual utility or risk tolerance. There are obviously different types of investors out there in the market who are willing to pay different amounts, and as Craig stated, the risk-neutral probabilities sum up the aggregate average risk aversion of all those market participants. Now what would Wall Street do? They normally would not pay for this risk. They would seek a buyer and a seller, charge their fee, and go on to the next deal.

In Table 1, I've just put these two models side-by-side for your comparison: the pricing model that uses what we call risk-neutral probabilities and the realistic model.

TABLE 1
BINOMIAL PRICING CALCULATIONS

MV	Outcomes	**Probabilities** "Pricing"	Realistic
$95	$110	25%	50%
	$90	75%	50%
		$95	$100

Here are some of the reasons why the risk-neutral probabilities don't necessarily equal the realistic probabilities:

- aggregation of utilities of all investors
- risk aversion
- uncertainty
- information asymmetry
- different models
- different model assumptions
- supply and demand

The main reason is that most investors in the marketplace are risk adverse, and the risk neutral probabilities reflect that.

Now, if we live in a world where most investors are risk adverse, so that the market value is priced below a realistic expected value for most investors, then we should expect to make a profit by buying a risky investment and holding it on our books. In the example that we just gave, the expected payoff was $5. That's in line with what everyone understands about risk and return. You take a risk, you expect premium in return for that.

Who in the marketplace makes a profit off these types of instruments? Now, we already mentioned that by taking risk you get that risk-premium that is priced into the market prices. For somebody who hedges the risk, if they hedge the risk entirely away, then they should pay that entire risk premium. They should not expect to make any profit. But there may be other sources of profit.

One is innovation. If you're first to market with some new idea, you can charge a premium above and beyond what would be the clearing price in an efficient market that's fully mature. The market maker, which is the game of Wall Street, is a type of innovator. The way they make their money is not by taking risk, but by coming up with new products that are out in front, and for which there's no actively traded market. They're able to get a large risk premium by doing that, and as that market matures and other competitors enter the marketplace, that risk premium will decline as that product becomes more of a commodity, and the markets become more deeply traded.

To wrap up everything that Craig and I have said so far, the essence of mathematical financial theory is that prices can be expressed as expected values. You do this via this tool called risk-neutral probabilities that really don't have any sort of intrinsic meaning. They're really just a mathematical trick to allow you to simplify the mathematics.

The concept that we expressed here in this very simple example with no time periods and no interest extends easily to multiple time periods. To do that you'll discount it using the short-term risk-free rate. If you make your intervals infinitely small, the concept extends to a continuous case.

What we end up with are the three major components of a pricing formula. You've got probabilities, cash flows, and discount rates. According to the Fundamental Theorem of Asset Pricing, if you have a complete model with risk-neutral probabilities, then you can take the actual cash flows under each scenario and discount them continuously at the on-the-run short-term risk-free rate to get your price.

OAVDE

$$OAVDE = \sum_{S=1}^{N} P_S^{\{\mathrm{Real}\}} \bullet \left\{ \sum_{m=1}^{M} \frac{DE_{s,m}}{\left(\prod_{j=0}^{m-1} (1 + I_{s,j})^{1/12} \right)} \right\}$$

(Note: Can be adjusted for Becker discounting.)

We have different probabilities. We have realistic probabilities. We have distributable earnings here, which are in fact the free cash flows for an insurance enterprise.

The other difference, other than the probabilities, is the rate at which you discount. You notice I did not label this as short-term risk-free rate. I just had a capital "I." I left it there because it's a bit ambiguous in practice what you're going to use for the "I." This definitely includes some provision for what the Treasury rates are along with this risk premium that Craig mentioned. This reflects the idea that there are different ways to incorporate risk into your valuation.
For market-traded assets, you can develop a model with risk-neutral probabilities, whereas in the case of realistic decision-making, actuaries tend to look at realistic

scenarios and express all of their risk in the discounting. It might be a cost of capital or target return. Luke Girard reconciled the two frameworks in some of his papers. He shows that you can come up with an interest rate that will equate the two.

To summarize these ideas, a risk-neutral framework is used for determining market values. What would you buy and trade these instruments at? Are you getting a good deal in the marketplace? The value of this computation is only of the mean. The individual present values under each scenario are meaningless in the risk-neutral framework.

A realistic framework is used for making a decision. What are you going to buy and hold long term? That's normally the question that we're asking in the insurance industry. Since you're using realistic probabilities of scenarios, your distribution of results, your present value of distributable earnings across a large set of scenarios, that distribution has meaning. You can express standard deviations and percentiles and statistics about the detail of the distribution, and they are meaningful in that context.

So far we've talked about interest rates, but I want to reiterate that the ideas that we talked about here apply to all pricing models as long as you have an appropriate model of all the relevant risks. As I mentioned earlier, for the types of products we normally deal with in insurance, that especially includes interest rates and equity performance.

Another big issue that comes up is that you need to have adequate sampling of your probability space. Here are the common stochastic variables that you might need to include in any given model:

- Treasury interest rates
- inflation
- equity returns
- equity volatility
- corporate bond spreads
- defaults
- foreign exchange rates
- foreign interest rates

I have eight things listed, but you don't want to have an eight-factor model. You want to focus for your particular application on what's relevant, on what you're valuing, and try to simplify your model down to the variables that are pertinent. Or to come up with a macroeconomic-type model where you can base all of these variables off some more fundamental economic variables to keep your number of factors down.

If you really wanted to apply the asset-pricing methodologies of Wall Street in insurance, you would need to have these things modeled stochastically, and develop risk-neutral probability measures for all these things as well. Now, in practice that's just not possible, but it is important that we adjust appropriately for the risk premium.

I'm going to talk about what the models look like in practice on Wall Street and in Insurance. For Wall Street, you typically need just a risk-neutral model, and what that involves is you need to have some projection of cash flows under all the scenarios. You need to have a good economic scenario generator with appropriate correlations between the various risk variables that are in that generator. You need hardware, which is one of the things that Wall Street has an advantage over insurance for sure. They have lots of hardware that is set up very efficiently to give them real-time analysis.

You need to have some sort of internal corporate process, a common understanding about what models are going to be used, and how information's going to flow. This process question is really a big one. It is something that Wall Street really has already invested a lot in, and has got things, obviously, working. This is something I think insurance companies need to spend more time working on, the education and participation of the relevant players in your company.

In practice, what ends up happening on Wall Street is, as much as possible, it tries to simplify its models down to as simple a model as it can use, and still state prices for a market. Then it normally talks of the prices in terms of some important input into that model, like volatility. Prices and volatilities in the option market are interchangeable information, that's what we call implied volatility.

On Wall Street models, in it's models, it is very interested in volatility because prices are changing all the time so the volatility it plugs into its model to get the market prices is going to change over time, and it changes very frequently. Wall Street maintains a very short-term focus on this, the volatility movements in the market. It's a little different than insurance.

This sums up the characteristics of a Wall Street pricing tool. The risk-neutral factors are as simple as possible to get the job done. You don't want to have 40 factors if two factors will do the job. Then it have this very large matrix of information that allows it to calibrate their simple model to actual prices for all the thousands of assets that are out there in the marketplace. There's always this tradeoff between being able to exactly price all the assets in the marketplace, and having a manageable model. You end up with a small number of factors along with this big matrix.

You can really think of the Wall Street models as a big interpolation tool. It knows what all the market prices are, and it came up with tools that allows it to quickly turn around real-time prices, and it keeps updating this matrix as it gets new information off the street.

Now, so far we've talked about what you need to model market prices. If you want to get into modeling realistic returns on some investment, there are a number of additional requirements. You need expertise in the underlying risk, as I mentioned before. If you're in the life insurance businesses, you need to have expertise in mortality rates. If you are selling catastrophic bonds, you need to know about earthquake risk.

You need to have a realistic economic scenario generator, and this is what makes the problem bigger for the insurance industry than it is for Wall Street. If you're trying to value distributable earnings on a block of life insurance business, then what that means is you're projecting those earnings along every scenario. In each point in time along ever scenario you need to emulate the financial activity of that company. You need to be buying assets. You need to be selling assets. You need to be valuing your liabilities according to the statutory reserve requirements, and setting up risk surplus.

To do asset purchases and asset sales requires that you do, in effect, a risk-neutral pricing at each node within your realistic model that you're using to come up with an OAVDE. At the same time, you need to go through the machinations to do a reserve computation at each node. It's a monumental task.

In addition, as I mentioned before, in the realistic model we have all these issues about what should the discount rate be. Once you have everything that we mentioned above, you need some way to weigh risk against return, and make decisions about what type of risk are you willing to take. You need some decision-making metric or methodology that weighs the average return that you get from computing an OAVDE to the risk, maybe the value at risk, or maybe the first percentile in the tail of your distribution.

If you've gone through the process of actually coming up with a utility curve, you can take all the results and then feed them through a utility function, and come up with a utility-weighted OAVDE in order to make your decisions.

The other thing that's different from Wall Street is that since we're taking this risky stuff and actually putting it on our books long term, the assumptions that go into our models are critical. To the extent that any one assumption makes a big difference in your results, you definitely want to know that. In fact, you want to know that up front when you're designing your product strategies, your investment strategies, whatever, so that you can try to design your products or your investment portfolio to be robust with respect to as many assumptions in the model as possible. Thorough sensitivity testing becomes crucial for actuarial work.

For life insurance, you have reserves. In Canada the valuation actuary subjectively sets reserves based on expert opinion, and I'm not sure how you deal with that issue. You get hedge accounting issues; you've got taxes and risk surplus. All these things you need to compute before you can actually get to the free cash flow of a life insurance venture.

Actuarial modeling tools need to be realistic. Instead of being simple they have a lot of assumptions: mortality assumptions, expense assumptions, and policyholder behavior assumptions such as lapses and elections.

There are run-time challenges, and to the extent that you're trying to value something that's not traded in a marketplace, you need to deal with the issue of illiquidity.

These are challenges that face both Wall Street and the insurance industry, but which are actually more challenging for insurance. One challenge is that to the extent that you come up with more complex products to model, you're going to need more factors in your model. This may lead to the need to develop some approximations to avoid excessive run times. Because of the issues involved in determining risk versus return based on not only the economic point of view but also GAAP earning volatility concerns, you've got a lot of business judgment whereas on Wall Street they can pretty much set up a mechanical process to do Delta or Greek hedging. For insurance there's a lot of expert opinion involved with making decisions. The other challenge is there's no way to come up with a closed form solution for these models, in the insurance industry in particular.

We've talked about the first two ways that Craig mentioned of representing risk. One is to represent it in the probabilities. One is to represent it in the discount rate, and what I'm going to suggest is that if we have any hope of integrating these two approaches, and I think we do have some hope of it, it'll help us make better decisions if we look at all these frameworks from an actuarial perspective.

But I think there's a third way, which in practice, is going to yield some good results in the intermediate term. The goal is to get as many of your risk variables into the probabilities as you can, i.e., develop economic scenario generators for all the relevant risks that are actually traded in the market, and for which you can get reliable prices. For everything else, what I'm going to suggest is that we try to tinker with the cash flows and leave the interest rate alone.

Now, we mentioned the goal of Wall Street is to market products, get an innovation premium, but mostly be just a market maker and collect a fee for doing that. An insurance company's goal is to take risk and to optimize its risk return tradeoff. It must earn some return, which means that it needs to get its risk up to the level that gets the target return. Within the universe of strategies that give it this expected return, it tries to minimize risk, especially the tail risk.

Another way to approach it would be to maximize expected profits given a target risk level. Now, it would be nice if we could use all these tools from Wall Street to value insurance because the risk-neutral probabilities already have some provision for risk, and you can measure a line's performance using risk-neutral probabilities to see whether they'd at least beat the market. If you can somehow rig your actuarial models for insurance valuation to compare as much as possible to prices on Wall Street, then you can do this comparison.

But the problem is the number of assumptions that are required in the financial theory, that are violated when it comes to insurance companies. One is that you don't have all your stochastic variables modeled. Normally you just have an interest rate, and equity return model, say, maybe an inflation component thrown in. But you don't have stochastic expenses. You don't have stochastic policyholder behavior. You don't have stochastic mortality or morbidity.

There's also the fact that the liability market is not actively traded, there aren't active prices, and they're, in fact, not fully marketable. As I mentioned earlier, we're always going to be using distributable earnings as the free cash flows when we're dealing with decision-making for insurance. I'm referring to this equation as the market value of distributable earnings, but I put a "T" here for theoretical market value of distributable earnings.

Line of Business Pricing Equation
(Discrete, Monthly)

$$(T)MVDE = \sum_{S=1}^{N} p_S^{\{R-N\}} \cdot \left\{ \sum_{m=1}^{M} \frac{DE_{s,m}^{\{ADJ\}}}{\left(\prod_{j=0}^{m-1} \left(1 + i_{s,j}^{STRF} \right)^{1/12} \right)} \right\}$$

This is a risk-neutral model for as many variables as have good market prices to calibrate against. You keep your discounting the same as it was in a risk-neutral model. Then you make all the other adjustments you need to make up here in the numerator. What that means is if you're not exactly sure what your expenses are going to be then you need to put some sort of risk premium into your expenses up here into the numerator.

If you don't know exactly what your mortality, future tax laws, or policyholder behavior, dynamic lapses, and dynamic elections are going to be, then you bump it up by some risk premium to account for that.

This is very much akin to what market makers do on Wall Street when they're coming up with new types of instruments. It has several advantages. One is you can use the risk-neutral models that already exist for traded securities. You discount at the short-term risk-free rate. Then you allow the actuaries to focus on what it is that they are experts in: policyholder behavior, mortality, and expenses. They should be able to put some sort of certain equivalent value on these instruments because to do so involves the same expertise as it takes to be willing to take this product to market and sell it in the first place.

Mr. Merrill: I have two comments. One, the suggestion that Russ had at the end is very similar to what David Babble and I have written about in a *North American Actuarial Journal* article a couple of years back. We talk about this approach, and we suggest exactly the same idea, the actuaries ought to do what they do best in projecting the cash flows. We deal with uncertainty in the actuarial risks in the cash flows in essentially an equivalent type of framework so that the financial risks that are market-traded and priced in the market can be dealt with using the risk-neutral probabilities, and the interest-rate modeling technologies that exist in the financial pricing area.

The other thing that I wanted to comment on is there actually is more connection between what I'd said in my presentation and the realistic probabilities that Russ had in his. Table 1 had two scenarios, cash flows of $110 and $90, with 75% probability of down, and 25% probability of up, based on the market prices. The way to think about that is a risk-adverse investor would not be happy with a 50-50 chance. A risk-adverse investor wants to be compensated. In order to take into account that need to be compensated, we adjust the probabilities to make the negative outcome more likely in order to rationalize the probabilities with the price.

Then there was the example with the 50-50 probabilities and an expected value of $100, and therefore, an expected return of 5%. That 5% is the risk premium that you have to add in when you're using realistic probabilities in order to do discounting. In fact, Russ, the realistic probabilities that you had are exactly the ones that I would be talking about in the equilibrium-modeling framework, and that 5% expected return that the risk-adverse investor desires to receive is the risk premium that you use when you're using realistic probabilities.

Now this is a key idea that comes out of the mathematics and is often missed in trying to get the intuition of risk-neutral pricing. If you have a model of risk aversion, so that you have some sort of an expected utility model for your representative investor in the marketplace, there is a one-to-one relationship between the risk-neutral probability and the realistic probability measure. We can go back and forth between those at will if we've adequately modeled the representative investor so that you can, in fact, use the risk-neutral pricing techniques in those actively traded parts of the financial market, and use realistic probabilities in those less actively traded parts of the financial market.

In fact, what we call the risk premium for those less actively traded parts of the financial market is an option-adjusted spread. When we talk about modeling the policyholder behavior, it's very analogous to modeling the mortgage holder behavior. The mortgage-backed security environment and the option-adjusted spread is the increment to the discount rate above and beyond the risk neutral probabilities for that unpriced portion of the risk.

There really is tighter link between realistic and risk-neutral probabilities than I think we sometimes appreciate.

Mr. Christopher J. Johnson: When looking at the liability cash flows, what do you think ought to be included in that? In particular, should a profit stream be explicitly included in the cash flows?

Mr. Osborn: Well, it depends on what you're trying to do. If you were pricing a product, I would say it should be distributable earnings that you're looking at. If you were trying to use these techniques for actually coming up with reserves, then you need to look at benefit payments, and perhaps expenses.

Mr. Johnson: How about for fair value of liabilities, should profits be included in that?

Mr. Osborn: No.

Mr. Michael S. McLaughlin: One of the things that FASB is doing for the fair value of liabilities is proposing a method that is similar to what's been discussed here where there are risk-neutral probabilities and projected cash flows discounted at a risk-free rate. One of the things that it's doing that's a little different than I think has been discussed here is proposing the use of the credit standing of the entity that is obligated to pay the liability. So one insurance company that's highly rated versus one that's not highly rated, all other things being equal, would come up with a different liability. I wonder what your views of that would be.

Mr. Merrill: I would say that this has a very limited chance of success in practice, at least in the near term. The problem is this: look at something as actively traded as the corporate bond market, where we've got credit ratings like AAA, AA, or A, and then you go out in practice and you look at AA bonds and you'll find that some AA bonds have spreads that overlap into AAA and A territory.

There is not a unique spread linked to each credit rating, and so it's a fuzzy characterization. It might be a step in the right direction, but with the limited trading activity on the liability side for the insurance companies and the limited data available for spreads relative to credit ratings, the chances of coming up with correct valuation I would say are pretty slim—at least until we find a way to gather better data on what the spreads are that attach to the proposed credit ratings.

Mr. McLaughlin: I agree there's a practical difficulty in coming up with what the spread should be, but I'm concerned whether it's even theoretically appropriate. In other words, there's the company assuming that it need not live up to its obligation. There's some kind of put option not only to shareholders if you're a public company, but also to contract holders, which I think is different.

Mr. Merrill: But, in fact, there is such a put option, and it's actually magnified in value by the guaranteed programs that are implemented to make the policyholder whole when the companies do fail to meet their obligations. Now, fortunately, insurance companies have shown a great fiduciary responsibility, but that put option does exist, and in fact, there are economically meaningful differences in credit worthiness in terms of ability to meet liability obligations across insurance companies.

Mr. Osborn: If you try to do this in the theory, there are a couple ways to do it. You can have scenarios where you don't actually pay or you could put an option-adjusted spread (OAS) in your discounting. I mean, if you go back to the Modigliani-Miller formulas that Girard uses you can get the same answer, or you can get an appropriate answer by putting in an OAS in the discount for the two companies, and it would be different for the two.

This was related to the problem that I often see come up in practice in actuarial questions is if you're comparing two strategies. You use the same discount rate for both prices, you may be missing the fact that the two strategies have very different risk profiles, and should require different discount rates.

Mr. David K. Sandberg: If you look at the role of the regulator, you would argue that they are maximizing their value by minimizing the value of the put option. In fact, the construction of risk-based capital framework over the insurance industry in the last 15 years has really been a way of trying to do that. I think we can expect the regulatory apparatus to do that. The Federal government has that same question, too. The banks went through the issue in the early 1990s, and people said regulators were failing in their job. So they revised the capital management process, or risk management process to minimize the value of that put option. The guarantee fund is there. I think in some ways it may draw more attention from a financial standpoint than might be warranted because, in theory, it should not be used if the regulators are doing their job appropriately.

Mr. Merrill: I agree with you. The guarantee fund is a Band-Aid on the possibility of the insurance company not meeting its obligation, and risk-based capital (RBC) is the Band-Aid on the Band-Aid because the first Band-Aid creates a perverse incentive. RBC is meant to remove the perversion from the incentive structure at a cost: profitability and opportunity cost for the companies.

Your comment also brings to mind the point that I wanted to reiterate. The financial reporting that does end up getting implemented has to make transparent what obligations the company has promised to pay, and what its ability to pay those obligations is. Those two things have got to be distinct and transparent so that regulators and investors are able to function effectively in this marketplace.

If you link the two together by projecting and discounting at risk adjusted rates, and hide the risk-taking behavior of one company *vis a vis* another, then I don't think our financial-reporting system is going to meet the needs of the marketplace.

Mr. Sandberg: I think that's at the heart of what Mike's comment was about. I think from the actuarial profession you don't want a new valuation process that says that as a company starts going down the tubes, its liabilities are going to reduce, and the market-value information that people are looking for to tell them about the company, is saying this company is still doing better than ever.

Mr. Merrill: If you allow them to bump up the discount rate for taking more risk and betting the bank, they make their situation look better in the financial statements.