RECORD, Volume 31, No. 2*

New Orleans Health/Pension Spring Meeting
June 15–17, 2005

Session 76PD
Stochastic Modeling in Health Insurance

Track: Health

Moderator: EDWARD B. MCELLIN

Panelists: JONATHAN MARK HENDRICKSON
EDWARD B. MCELLIN
ARMAND M. YAMBAO

Summary: This session covers stochastic modeling processes currently in use and explores others that may be beneficial to health actuaries. Some areas currently impacted by stochastic modeling include reinsurance, capital and surplus, anticipated earnings, provider contracting and claim liability estimation.

MR. EDWARD B. MCELLIN: I'll be the moderator for this session. I'll go through my bio a little bit and then let the other presenters do their bios. I'm from the University of Melbourne in Melbourne, Australia. I'm from the United States. I used to work for Blue Cross/Blue Shield of Florida and then decided to try something new and do a little bit of academia. The center is a joint venture between the Institute of Actuaries of Australia and the university and also some private businesses in Australia. I hopefully will present some practical methods and not just academic methods. If you downloaded the notes, you know there are some summations and things like that you might see from a Bower's book, which is scary, but hopefully it won't be that bad. People see the stochastic modeling, and it's the same thing with me, too. I hardly used it at all, and I now use it all the time for businesses there.

Mainly I'm going to be presenting, introducing the topic and comparing it to some of the deterministic and statistical stuff you use, as we all do, and it's easy to understand. Then I'll give you an easy example for incurred but not reported (IBNR) claims, which most disability and health insurance people know all about, and property and casualty (P&C) also. Then I'll compare both methods so you can see

Note: The chart(s) referred to in the text can be downloaded at: http://handouts.soa.org/conted/cearchive/neworleans-june05/076_bk.pdf.

how if it's a stochastic method, you can do the method, figure out the variables and look at that as a better method than some of the old statistical methods. There's cost-benefit analysis, but we'll get into that. I'll let the other presenters introduce themselves when they come up. We'll have Armand next. He'll be giving you more practical examples of how to use some stochastic modeling in health insurance.

Let me go through the generic right here, the things you've seen when you've done professional development and Course 7. We'll look at what we might want to investigate and then cover the four things I've always been concerned about, such as consistency. The model's going to have pretty predictable results, showing what you want to forecast. There's also realism. You have to have the results you think might happen. The model also has to be results-oriented and accurate. A lot of these models can be checked later on to see whether they're accurate. I'm sure you've gone through and tweaked a little bit to make a more accurate model in the end. There are other models, but I'm just cautioned about deterministic statistical models and stochastic models. A lot of these can be more individual and collective models, but most people use deterministic models, and I don't think I used stochastic models before I got out of Blue Cross that often. I knew what they were because I learned them in college, but then I started using them more in academia and then put them to practical use.

This session is about stochastic modeling, so I don't want to go too far in deterministic, but I want you to get an idea of the differences. This is the main thing. All the variables are known in advance. You know everything about a deterministic model. Everything's fixed at the outset. There's some formula. You put in your data, and it's going to turn out a point estimate. You'll see that example easily with the IBNR. It's going to be an actual number; there's not going to be a range of numbers. If the formula's the same and the data are the same going in, you'll never have a different result coming out. You have to change the assumptions or change the formula to get anything different.

People understand this, and I do, too, and that's why I've always used it. It's how things have been done in the past because we didn't have computers to run thousands of simulations, so it was easy to understand. All you had to know was the formula, and you could even use it before spreadsheets. You put the model in, you put the data in, and the formula came out. It's easy to understand. If you change one parameter or more, or if you change the formula around, you know exactly how it's going to affect it, especially if you graph the results. That was a nice thing.

The bad thing, and that's the main thing with stochastic models, is you don't know the probability of an event happening. In my example, we'll put in an IBNR number of 10,000. Chances are it's not going to be 10,000 for a certain period. There has to be a probability around there. What's the outcome going to be? That's the thing with stochastic modeling. You know the probabilities, and everybody knows the normal distribution, so you know 99 percent that something is going to be within three standard deviations of the means. That's an easy distribution. There are a lot tougher distributions out there, and when we're doing the example I'm talking about, we're not going to be using an easy, normal distribution. We'll be using other ones, but they're all based on the same kind of thing where it's taking on a distribution that patterns the behavior you're looking at, and that will be on some of the other examples, too.

I'll give you an example. If you're using inflation or pricing (we always did that in our exams), something you said was 2 percent, that's a deterministic. You're just saying it's 2 percent, You churn out what the premium's going to be in 10 or 20 years. That's a deterministic model, but you can always model inflation or something else, mortality or anything, by a random variable. What's inflation going to be? Some of the stochastic models are a little bit better because they imply some of the randomness in there.

Talking about this, the main things are the random variation and the variables, such as mortality, inflation or interest rates. It's unrealistic to think things are going to be the same 20 years from now. Maybe on average they will be, and in a lot of these models you do, you can get a good result. You can do an average of what something's going to be, but it's best to have some distribution for some of these variables and go through there, and you'll get some probability behind what's going to happen. We still have to figure out some of these distributions, and we don't know what those are going to be. Once you see some of these other notes later, I can say that it's nice that we get all these things, but we're still making guesses and assumptions about distributions and what the parameters are going to be.

We all know that interest rates vary, so if you use the old methods, such as Bower's, a fixed interest rate is not realistic in the real world. You want to have a randomness. You also want to know probability. I'll give you the results. I'm 80 percent sure it's going to be within a couple standard deviations of that result, and the main thing now is we can model these. For a lot of these things, you can use computers now. They couldn't run 10,000 simulations in the 1950s and the 1960s. You keep running these simulations now and get a little bit of a pattern of what's going on. You still have to know the distribution and what the pattern is. You can't use a normal distribution if something's lognormal, and that's what we'll get into.

I want to go through a quick IBNR. We're comparing some deterministic and stochastic methods. I hope everybody knows IBNR from disability insurance and health insurance. It's just reserving for a claim that you know has been incurred in the past but hasn't been reported to your claims processing yet. It's used in health, disability and P&C insurance and implies the time lag before the claim is reported and settled. There's some liability that we have to account for. A lot of claims aren't paid, and most of the claims processing departments do their accounting after the month or after the quarter, but we have to know what those unpaid liabilities are because that's something we have to reserve for.

I've used this at Blue Cross, where we tried to set something to a loss ratio or tried to look at the past and see what it's been by the deterministic or statistical methods. I'll use the terms interchangeably. I've done it both ways where I've looked at three years back and seen that something that passes that way should be the same in the future, but we know that's not true. Or we set it to some loss ratio, such as 86 percent. I'm going to try to go through this quickly because you're not trying to learn deterministic methods here; you're trying to learn some stochastic. Everyone knows the chain letter method. That's a statistical method, and it's still a good method for IBNR, but we're trying to expand on that.

Traditional runoff method is based on the concept of estimating future claims. You've seen these triangles. You might see the lower triangle, upper triangle. You're trying to determine everything below that triangle, filling out something that occurred and how it's going to run out (see McEllin Slide 6). I've only given a small portion of it here, but this is all cumulative. March was the incurred month, the paid month, and then how much is going to be paid down the line. Certain types of businesses have a longer runoff than others. This is not a bad method, but it's just going to give you some kind of point estimate in the end, and that's one of the main things. We're trying to get a probability of what's going on.

One method I'll go through quickly is completion factors, or development factors, which I call them in Australia, just using past data. I won't get through a lot of these variables because that's when people tune out. You'll see a lot more of that with the stochastic modeling. You're not going to learn how to do stochastic modeling of IBNR from this, but I'll send you the paper for it if you want. Hopefully you'll get a good gist of it here. I did this when I was in my first year at Blue Cross.

McEllin Slide 7 is just looking at the past, dividing into that how many things completed out by looking at the summation of your IBNR and then going forward and seeing what you have in the current month. You can see that generally it runs out in 10 periods. That's how it's going to be later on. You get a point estimate of the completion factor going forward and apply that to other methods. That method is okay if the liabilities aren't too big, and you don't care about having a probability distribution around that IBNR. I've used it many times myself and had a pretty good accounting of it. I didn't want to stay that much on the deterministic method. Like I said, there are other methods, and I compared four or five methods that I can't get into now but can show you if you want to know a loss ratio method, a statistical method and stochastic method and how to determine IBNR.

Let me go quickly through the stochastic. Stochastic IBNR is retrospective in nature. It's not stock prices, so you're a little bit better, where the past doesn't determine the future. especially if things change in your long runout periods, so we never want to rely too much on the past. You'll see even with stochastic we're going to have to rely a little bit on the past, but not as much as with deterministic, statistical methods. That's all we're dealing with with deterministic. We're looking at the past. If things are like that, if they ran out within four or five periods, it's going to run out again in four or five periods, so we'll apply those completion factors and get a good result for that estimated IBNR. We're trying to have a little less reliance on past data, and we're still going to have to rely a little bit on that, but that's one of the reasons for a statistical model. I just want to measure, to see, what the probability is. I don't want a point estimate. I want to know whether something's going to be 20 percent within the reserves I'm calculating or 10 percent, Mostly it's distribution we're working with. If we can get that, it gives the higher-ups a little more confidence where our reserves are going to be and what we're saying they are going to be within the 10 percent guess. That's one of the good things about stochastic methods.

This is just a whole thing in all stochastic modeling. This is IBNR, but it can be other examples, too. We'll go through the variables that affect the outcome. That's easy. What are the probability distributions? That's a little bit tougher, but a lot of research has been done on this, and you can do your own research to figure it out. The parameters are straightforward, running some matching estimates or things like that, and then printing all these simulations that will generate some of these out. I'm not going to give you results here in this presentation, but I'll give them if you want afterward.

Let's do it for the IBNR valuations (see McEllin Slides 9 and 10). This uses random variables and their corresponding probability distributions. We have every incurred period, which is one of the good things about stochastic. We're looking at every period. We're not summing everything together. We're looking at everything individually. We have to know how many claims we had in the period, the severity, how big the claims were and the claim amount. Remember the N and the X. Those are the variables we're going to talk about. The report lag is the main thing. The other things can be with any kind of calculation for claim amount, and again you're seeing these in a lot of books and Bower's. This is the intriguing thing. What's the time lag between an incurred date and the date it's reported? The other things have been written down for years now: how much they claim in insurance claims or health claims or how many claims are going to be there. We're concerned with the report lag because that's what's going to determine our IBNR.

I'll go through this quickly. Here are a lot of summations. Again we're defining the variables. Let T be the time of the claim. We've always got to assume where it's going to happen. We'll make it easy in the middle of the period. That's the big indicator there, that third bullet, whether the claim was reported or not, and then we know it's an IBNR claim because we can look at these things and know whether it was reported later than the month it was incurred or any period. You can look at it in years or quarters, whatever period you're defining. Y is the big IBNR. We have X, which is the claim amount, and N, how many claims they're worth. We're summing all the claims over. I is that indicator variable. If the lag is bigger than when it was incurred, we put it into our database, and it's going to be affecting IBNR. If it's not, if it was paid the same month it was incurred, it's thrown out.

The problem's pretty simple, but what the distributions are in here is a little bit more difficult. Again you're seeing all these summations that could be scary, but when you put them in a spreadsheet (or a lot of software does it for you), it makes it a little bit easier. Total IBNR is every individual period we're summing up (see McEllin Slide 11). I'll be talking about one of the big assumptions that we have to apply here, or all these will start being thrown out of whack. That's independence. That first line is total IBNR, but if we can assume independence, none of these things—the number of claims, the claim severity, where the probability was and whether the claim was an IBNR claim or not—are related to each other. Those things make stochastic modeling easier when you can assume independence. If we didn't assume independence, if one thing was affecting the other, we'd have a big experience matrix, and that would make it a lot more difficult, so you make all these assumptions that make it a little bit easier.

I left out other things that could be affecting IBNR. There's settlement claim, where you know the claim has been reported, but it takes three or four months for the claim to be paid off, I could have added that variable in there, too, or any other variable you think might make it long for a claim once it's reported or once it might be incurred to be paid off. We just added this independent on to the end. You need to know distribution for it, but it doesn't make it that difficult. Once you lose independence, that becomes a big, bad monster to try to figure out. Those are the things you need to make sure your assumptions are right. With independence, we can easily make our expected claims each individual function now.

You're not going to learn how to evaluate these things at this period, but it should hopefully get your mind set as to how things are going. We have that formula for IBNR. It is just how many claims. We have the distribution for that, or we will have in a little bit. What's the claim severity and what's the time lag? We need to know all the probability distributions for these things. What's the probability distribution? You can look at the past data, which you're relying on. It's a little bit better than with the deterministic methods, but look in the past to try to figure out the general amount. What's the amount of claims? Do we have 10,000 claims every quarter? Is the amount usually $2,000 to $10,000? We could look at past data to see how these distributions are, to make it more personal by the company.

Generally we have information on what a lot of these distributions already look like. You can look at any actuarial science book and see the distribution for claim amount and number of claims. Distribution claims are usually poisson and negative binomial. Poisson is preferred. I didn't put the points on distribution on here, but mean and variance are the same. You estimate only one parameter. That's the next step in this, and that's a little bit difficult, too, and it's additive. We don't have to worry. We just keep adding these things up and get our claim frequency and the total IBNR, so that's a nice quality of poisson.

The thing about the negative binomial is it doesn't have an additive property, so mean and variance are different, and that makes it a little more difficult. You want to make finding these parameters a lot easier, and if you do poisson, all you need to know is lambda. You can look in Bower's. I don't want to go through all these things, because I hated Bower's myself, but all these things are your distribution of some of these things.

There's no theoretical support. You can look in a lot of the papers about what the claim amount of things is, but there are certain things that we use on claim amounts. We know they're one-sided, meaning there's no negative claim field. They're highly skewed. They're not a normal distribution. There are some big ones up there. You have some centered at $2,000, but you have some huge claims out there for $50,000, so they're not a normal distribution. There are some popular values that they're centered at. If you know these properties, you can look at distributions out there, and people have done this. Gamma, lognormal and Pareto distributions are considered. This is in all actuarial books.

This is a thing that people don't know about, and this is where it's going to determine your IBNR. What's the report lag? You can examine past experience. You can see how long it takes for one claim to be processed before another claim, and that's waiting time that I've looked at a little bit. What's the waiting time before the next event occurred? We've all had queuing theory, and waiting time for one claim is why we have queuing theory—when the next event's processed. That's a distribution, and we have to assume some distribution and mostly waiting time is exponential distribution. There are some problems on there trying to determine what the queuing time is, and it is the same time with the lag between something processing through the system. Exponential, poisson and lognormal are used for distribution. We generally use exponential distribution to model waiting times, and we can use that now for report lag, but you can make your own distribution. You can use the past data. That's using some of the past data or not relying on it to the extent we do with the deterministic methods.

I'm going to discuss generating results quickly. We have three random variables. We can have more. I took out settlement lag and other lags. That shows you another distribution of what might contribute to something being paid on IBNR. We have number of claims, claim size and report lag. Number of claims and claim size have been used often about distribution. It's the report lag that might be different between different companies, and so give each an appropriate distribution.

Here are the parameters. All these distributions have parameters. Poisson has the lambda, and there's exponential distribution, and we have to find out something about what these distributions are. That's going to be different for each insurer. That's again where you're going to have to use some of your past experience. What's the mean? What's the spread of some of these things? That's where you get into stochastic modeling and running all these simulations. You still have to have a good basis on what these distributions are. I wish I could get into that a little bit more. I'd be happy to talk to you a little about it because that's applying this practically. How do I do this? All these distributions I talked about can be run on

Excel, which is simple, but there are a lot more powerful simulators and Monte Carlo programs that will run all these simulations for you. Those are things that you'll have to get into more, not in a 20- or 30-minute discussion.

Optimal parameters can be found by using maximum likelihood, moment generation and iteration. What are the best parameters for these things? Iteration's easy. I saw what was in the past. I ran all these simulations to see what's close to what I saw in the past. I'm relying on those past data again, but again I'm trying to find some of these parameters that I need for these distributions and then running all these distributions again and getting 10,000 trials and seeing some probability of what's going on. This is the gist of it, which I'm only scanning over, but I think you'll see when we do some more examples that will apply to some things how you might not be not seeing any results here.

Why do a stochastic over a deterministic or statistical model? We're looking at each claim individually. For every trial, we look at each claim. We're doing some probability for that. The number of claims is going to be this, the claim severity is going to be this, the lag is going to be that, and we're running these trials for each of them. It's based on experience, but if I put in a normal distribution for the number of claims, it's going to be completely off. You have to get these things right. You have to determine the distribution and the parameters. The distributions are important. What does it look like? You can look at your own data and see what it looks like. Does it look like a poisson, lognormal or exponential?

You don't want to rely on too many data, but we have to. You can learn something from the past. You don't want to rely entirely on the statistical method, so again parameters need to be estimated from the past, unless you have a good actuarial knowledge of the mean, the average claim that you're seeing over the years or the number of claims. You can stick those in easily.

The main thing about these stochastic models is it's probability, but I'll give you an IBNR amount and be 90 percent sure based on my distributions that that's going to happen, and then we can easily check that in the past. I just skirted over and didn't show the results of some of these things, but if you want to, I can show you the paper that compares all these methods: the statistical method, the stochastic method, the loss-ratio method and some more involved statistical methods than stochastic. A lot of them have similar results; a lot of them are different based on garbage-in, garbage-out assumptions. Don't be afraid to use some of these stochastic methods because they're not that difficult with computing powers. Excel will even do a good job, but I wanted to lay the foundation on some of the stochastic modeling.

We're going to get into examples from some of the other presenters showing you the results on some of their modeling with health insurance and then at the end we'll take questions. If you need more on that, I can always give you a paper that compares the methods and runs the insurance through some of the actual Visual

Basic and Excel.

MR. ARMAND M. YAMBAO: I'll describe my background as far as my work and give a picture as to what kinds of examples I'll be showing you here. I work for Ennis Knupp & Associates as the manager of financial modeling in Chicago, and we are primarily an investment-consulting firm. We do give asset allocation advice, for example, to funds that have health-care liabilities attached to them and sometimes even pension liabilities attached to them. The two examples I will describe are the latest section of my presentation. One's for a health-care fund for a public sector. It's for a group of teachers, who set up a new retiree health-care fund, so the fund has a long time horizon. We need to advise it in asset allocation appropriate for its type of liability and the retiree health-care benefits that it is promising.

The other example that I'll provide is for a corporate plan, which also sponsors a retiree medical benefit. In this case the context is different, in particular the regulations governing the plan's retiree health care. Corporate plans are different in terms of accounting and how much expense they have to report compared to a public fund. This will give you a variety—the same kind of problems. We have to determine the right asset allocation for retiree health-care funds, but in two different contents.

The purpose of the stochastic modeling exercise that I've shown here is to try to find a good strategic asset allocation. As you will see later in this presentation, it's not like once you've done it, there's some magical answer such as, "Go 50/50 equity or fixed income." It's more of a tool for the fiduciaries to guide their decisions. If they go to this much aggregate, for example, what does that mean in terms of risk for them? It translates the concept of investment risk and is something they can readily understand in terms of their cost of maintaining the plan in the long term. That's the major reason for strategic asset allocation. It's more of a guide or a helper tool to do that. In this type of exercise, when you're dealing with long-term projections, the advantage of stochastic modeling is paramount. There's a lot of uncertainty for these types of liabilities or projections.

I'll go over our methodology and how we develop our assumptions. Basically all of these are infinite time horizon, long-term, until people retire. We looked at a 15- or 20-year time frame, especially for corporate plans. They're concerned with their cash accounting expense over the next five or 10 years. We still take into account the long-term liability at the end of that period to capture the long-term strategic nature of the plan, but the timing horizon of our projection is usually within 10 to 15 years. To do that, we pretend that we will do the valuation of the valuation of the liability over that 10- or 15-year period and do it 500 times, 500 files or 500 scenarios to capture the range of uncertainty on those future valuations.

As Ed mentioned, the first step in any stochastic model is to find out what the variables are. In this case, because we're setting asset allocation, the most important thing that you want to model correctly will be the economic factors or the capital market assumptions, The most important thing when you do this type of asset allocation study where you're trying to set an asset policy into your liability is to make sure that this set of economic variables that we're trying to set is valuated in a consistent manner, in the way that you're using those economic variables for your liability.

In this type of case, the most obvious economic variable that ties everything together will be inflation, and more important, health-care inflation and how that relates with the way we will simulate asset returns and yield rates to value the liabilities going forward. Here, unfortunately, you can't assume independence between the variables because it's not a realistic assumption when economic outcomes transpire over the next 10 or 15 years. In reality if you have a high inflationary environment, it's more likely that interest rates are high. There are some extraordinary events that might lead to some inconsistencies, but you want most of your trials to have some consistency in that regard—the high inflationary environment will lead to higher interest rates and high yield rates, and that will constantly affect the fixed-income returns. From there on, it will also affect the equity returns to the extent that you will model the average equity risk premium on top of fixed-income returns.

To handle the complexity of the dependence of these variables, we start off simulating each random variable or economic variable one at a time, and you do the next one dependent on what you have simulated for the previous one. What we do is, although we say that is the random run or simulation, to some extent the later variables are not completely random anymore because they're now dependent or the parameters will be restricted to what you predicted the other variables to be. We start with the basic one, inflation. Another important thing about starting with inflation is that this will make sure that the inputs for your liability will be consistent your input for your asset projection. For any type of simulation model, and here we have dependence, it's important that you choose the right one that you need, a building block method, where you start with and build off your other simulation outcomes.

For price inflation we assume an average of 2.5 percent going forward, and we anchored that in the expectation for inflation with a blue chip economic survey forecast by expert economists. We build out variability around it using historical variability and also treat inflation as having some mean-reverting factor and taking into account that Feds will come in to regulate the policy to insure that if you have a sudden spike or inflation shock, Fed policy will make sure that eventually the mean will revert to our long-term historical target, in this case 2.5 percent.

After building the basic price inflation, the next important one will be health-care inflation or the rise in the benefits that were promised for retirees. We start with the price inflation that you projected already and then build on some spread to it, looking at historical information in terms of covariance of the health-care inflation versus that baseline general inflation. We would generally assume an ultimate period where over the next five years we'll have slightly higher inflation until it tapers down to somewhere closer to the general-price inflation.

The next variable developed would be 10-year Treasury, or the yield rates going forward. We just used the 10-year Treasury as the midpoint, somewhere in the middle of the yield curve, but we also simulate the other parts of the yield curve. Given that current yields are at their historically low level, we assume again a mean-reverting trend that it will slightly increase over the next five years and level off from there. Again building off some variabilities around it, but to develop that simulation, we have to assess the average of the outcome after we do that. That's the target that we set. It will revert to some expectation within the next five years, something lower compared to historical levels, but a rise in trend.

We saw the inputs for the liability side of the equation. This will drive how the liabilities that we're looking at evolve over the next 10 or 15 years. Now, what are the variables that will drive the asset projections, again making sure that we anchor what is done previously? Fixed-income returns are related to changes in yield rates. For example, a rise in yield rates will result in lower fixed-income returns. Overall, looking over a 15-year period, our average fixed-income return expectation in our model is about 5.8 percent, with lower returns during those first five years where we expect a steady rise in yield rates until it levels off after that.

For U.S. stock returns, we use a building block approach where we view stock returns as being driven by inflation, and the premium on top of it is how much it will earn in terms of dividend yields from those stocks that you own and how much we expect in real growth from the capital that you have invested for those stocks that you have owned. We have a consistent baseline for inflation, and right now the current dividend yield if you look at the Wilshire 5000 Index is about 1.6 percent (it's slightly higher right now), and our expectation for real growth in earnings is about 4.7 percent. That gives us the total expectation of our U.S. stock returns to be about 8.8 percent over that 15-year period. Again, I mentioned here, we have a slightly lower return.

In the capital asset pricing model, we look at historical data or historical market risk premium for other asset classes such as real estate (for example, the fund wants to invest part of its money in real estate). If the fund wants to invest in non-U.S. stocks, we do the simulation also using capital asset pricing models. For covariance tests, we use historical estimates. We believe that they're more stable.

Before we go into the first example, to give you a feel, Yambao Slide 8 shows the range of simulated asset returns compounded over a 15-year period. The horizontal axis represents the amount of equity, or stocks, that are invested in the different scenarios that we'll be looking at. It could be as low as 5 percent in equity, meaning 95 percent in fixed income, or as much as 100 percent in equity and nothing in fixed income. As you would expect, investing more in equity will give you a higher return (that red dot in the middle is the mean of the 500 trials). This is a floating chart, represented by the 5th and 95th percentiles. As you go to higher equity allocation, you have a higher average return, but your 5th and 95th percentiles represent the great uncertainty of your long-term return if you invest in more equity.

If the fiduciaries didn't do the asset liability exercise that we will show later, just looking at this, it's hard to translate the risks and rewards of investing in more equity. This is the asset return, but what does it mean for them in terms of costs for the plan? It's what we try to do when we do the stochastic model.

The first example that I'll discuss will be for the public fund. The horizon is much longer, looking at a 50-year time period, so we extend the projection to 2052. Before going to the stochastic model, here's an example of a regular deterministic model (see Yambao Slide 9) where we say that we'll change the asset return assumption only, and we'll change it only three times, whether it's a constant 5 percent, 7 percent or 9 percent. We'll concentrate on the closed group, meaning there'll be no new entrants coming into the plan going forward. More realistic is an open-group plan, where we assume that new participants will come in, and the population will grow. As expected, more retirees will come from the system later. That's the reason why in an open-group projection, you ran out of assets within this 50-year period: you have more retirees coming out of the system.

Given that this is a public fund, one of the constraints is that it can't increase its contribution too much. Even if there are a lot of active participants coming in, contributions are fixed as a percentage of payroll such that a lot of the costs are borne or are put into the strain of the assets, so that within that 50-year time frame, the fund ran out of assets given the expectations of health-care inflation.

To translate the second graph in this deterministic model compared to a stochastic model, what we're seeing now is that instead of just three lines representing 5 percent, 7 percent or 9 percent asset return, we're seeing a range of the 500 trials that we've done in the horizontal axis, the time period from right now until 50 years from now (see Yambao Slide 10). You'll see that the red dot in this case (the average or median) follows the same trend as the open-group projection, the middle one where the assets run out in summer 40 years from now. The good thing about the simulation approach, or the stochastic model, is that you see the full range of uncertainty. As you would expect, going out much further in the future, you have greater uncertainty. This is assuming the baseline of 60 percent risky asset allocation.

The next step is to change the asset allocation target. How would that help? We try to capture that by looking at the risk/reward chart (Yambao Slide 11) where we say the benchmark is to stay where you are right now, currently invested at 60 percent in equity or risky assets. Each dot in the chart, for example the 70 percent, represents the change in long-term economic costs if the fund increases its equity to 70 percent. What we did for a reward measure is determine the change in costs for moving to a higher equity allocation. The vertical axis is the reward measure. You will see about a $500,000 long-term cost savings on average if you go to higher equity. It's intuitive that moving to higher equity offers long-term savings, but what's the corresponding risk?

The corresponding risk is determined by figure out what the cost increase could be if instead of having an average economic outcome we have an unfavorable economic outcome. What's the downside cost increase? To determine that, we need to narrow down the average. Instead of averaging all 500 trials, we average the worst 100 trials. Out of the worst 20 percent of outcomes, what will the average change in cost be? This time, by going to higher equity, instead of seeing a cost savings, you see a cost increase. That's going to the right of the horizontal axis, where by going to a 70 percent equity allocation, you could potentially increase your long-term costs by slightly less than $500,000.

Is that an acceptable risk for the fiduciaries to take? Going to higher equity can increase your reward in terms of lower costs for the plan. Show me again that average economic outcomes happen. If there's an unfavorable outcome, is that an acceptable tradeoff? We plot all the other scenarios, all the other ranges of mixes and risky asset allocation, ranging from as low as 25 percent and risky to 100 percent.

Our results will not jump out with an answer that the fund should go 70 percent or go 50 percent, but they provide a better understanding for the fiduciaries to sense what it means when they change their asset allocation. It's more than saying, "You'll earn 7.5 percent instead of 7 percent." We're saying, "Your costs will decrease by this much on average, but there's still a downside risk that you could have higher costs." The question is, Where's the comfort level?

As a general rule, it's hard even with this information for fiduciaries to make that decision. At a minimum, what's helpful in the chart (Yambao Slide 12) is that there’s this line saying that if you go past 80 percent, your incremental reward or cost savings is starting to be less for each incremental unit of risk that you're taking. A mathematical way to say it is that your slope is decreasing, and that's probably a good stopping point for taking on risk. You don't want to go beyond 80 percent, for example, in equity.

In the same way you're going down the curve, your potential cost increase by going too low in equity is so much higher for each unit of cost savings that you're seeing in a downside economic scenario. You probably don't want to go way down below in that risk spectrum.

Other measures of success or failure, or risk and reward, could be when your assets will run out, whether you're going for higher equity allocation, whether your reward will stay there for a longer term or whether you'll run out of assets in 2050. You run the risk of increasing your chances of running out of assets within a shorter time period, meaning 2030. The way we look at this is where's the change in slope? Is there a significant change in the chance of running out of assets by going from one asset allocation to the other?

The next example (Yambao Slide 13) is for a corporate plan. This plan doesn't have a long time horizon. It's only looking at the next 15 years. In a stochastic model, this is what the expense looks for the retiree medical plan. In this case, before looking at the long-term economic cost measures, we'll look at the surplus return efficient frontier (Yambao Slide 14). We define surplus return as the liability growth for the health-care benefits that are promised less the asset growth. The reason we went to this exercise first is to look at different asset classes in particular, not just risky versus nonrisky. It's a little bit complicated, where the plan wanted to look at fixed income, Treasury inflation-protected securities (TIPS), small cap and large cap stocks, non-U.S. stocks and municipal bonds. We tried to narrow down the choices first by looking at the surplus return efficient frontier.

If you remember Course 6 of your investment exam at the SOA, you heard about the asset return efficient frontier. What this does is, instead of just looking at the asset return efficient frontier, or the reward and expected return versus the standard deviation, we're looking at the expected surplus return and its standard deviation given different asset allocations. What's worth mentioning is that the red line on this part is the result from an asset-only efficient frontier, and the blue line is from the surplus return efficient frontier, meaning that with the low-risk spectrum, you see some differentiation in the results. What happens is that under surplus return frontier analysis, you take into account the characteristics of your liability, meaning you should probably consider more TIPS than regular nominal bonds if you're trying to match your liabilities with your assets under a low-risk profile. Our main message here is to confirm what you intuitively expect the model to generate.

Looking at the cost of the plan in terms of long-term costs, here's the same type of chart I described for the public fund example (Yambao Slide 15). There's a different rule for calculating expense, but there are the same basic results. The reason I have two 100 percents is, from the previous efficient frontier graph, I've chosen nine asset policies, or nine asset allocations, and the last two, the riskiest ones, both are 100 percent in equity-type investments. It's a different allocation, whether you have more in small cap or large cap stocks. These are the nine points that I've plotted on this risk-reward chart. There's the same kind of message: What's the right risk tolerance for the fiduciaries?

MR. JONATHAN MARK HENDRICKSON: I'm an actuary with Milliman in Phoenix. We're going to talk briefly about the central limit. We use it in Monte Carlo simulation, but you don't have to worry too much about it. I'm going to go into Monte Carlo simulation and give you a practical example of the simulation and go into some fundamental questions about how you would set up simulations. We'll then go into a particular example that can take a full-risk situation for a health plan or a large group into account. We'll talk a little about specific reinsurance and show you some results based on that. We may have time to talk about aggregate reinsurance and Medicare Part D, or we may not.

First of all, what does the central limit theorem say? For a given distribution with a mean and variance, the sampling distribution of the mean approaches a normal distribution with mean μ and variance σ^2/N as the sample size increases. The key is that it's a distribution in and of itself of the sample means. This works for any distribution, as long as the variance is a finite number. We're going to get into an example where we start with a nonnormal distribution and by taking samples of it, looking at the averages and looking at how those averages are distributed, they turn into a normal distribution.

The key, in my opinion, is that N is the sample size for the sample, not the number of iterations. You can have a sample size of five, and if you can run it 10 bazillion times, and it's still not going to be normal depending on the underlying distribution. That's the key aspect of that. N is the number in your sample.

How does that impact Monte Carlo simulation? We're going to talk a little about it. You start with the probability density function (PDF). What's the most common PDF as health actuaries if you do traditional health work? That is a PDF. We convert it to a cumulative distribution function, and we can generate some random numbers and pick up outcomes from that PDF. You do that by calculating the inverse, which you can do easily in Excel, and repeat that N times in each sample that you want to do this for.

The fundamental questions that surround your Monte Carlo simulation are the number of iterations you perform and the number in your sample. The number of iterations is subjective and depends on how smooth you want your results to be and what your underlying distribution looks like. You should consider any computational time restraints that you may have. In the example that I'll show, I started with 50 lives and ran up to 10,000 lives and did 1,000 iterations and 10,000 iterations. I can run it in about 15 minutes. The increases in the laptop and desktop speeds have helped this process. When I first did this 10 years ago for a reinsurer, it took considerably longer. I had to work all day setting it up, started running it before I left, and it was still running when I got in the next morning.

The design of the spreadsheet that you're going to use to do this is a key component. If you go out on Yahoo or Google and search for Monte Carlo simulators, there are a couple of popular ones out there. Most of them use Excel as a backbone to what they do. Some of them don't use Excel as a backend, but some do.

The number in the sample is what we're going to look at results based on, so we're going to show you results for 50, 100, 1,000, 5,000 or 10,000 lives. The underlying distribution, in our case, the CPD, is going to determine how close to normal the averages of our samples look.

The next two slides will illustrate both fundamental questions. How many iterations should you do? What's the effect of sample size? I'll show you the CPD I used in a second, but I'm going to flip back and forth between these slides. Here are the results of 1,000 iterations (see Hendrickson Slide 2, page 3). You can see it's jagged but the 5,000 and 10, 000 life samples look normal. Let's switch to the 10,000 lives (see Hendrickson Slide 1, page 4). It's smoother, but there still are some jagged edges, especially for the 50 and 100 life groups. What strikes me about those two is that they don't look normal to me. I'm not sure what they look like, but they're not normal. Even for the 500 life group, the tail is longer toward the end, but when you hit 1,000 lives, it's starting to look normal.

What I did is take a medical CPD as a PDF. The underlying average cost of the CPD was $3,000 annual claim cost ($250 per month per member (PMPM)). In this example, I'm also going to run with specific reinsurance amounts of $25,000, $50,000, $75,000 and $100,000. The group sizes are 50, 100, 500, 1,000, 2,500, 5,000 and 10,000. I decided to go with 10,000 iterations.

Here are the PDF and cumulative function (see Hendrickson Slide 1, page 5). Usually when we look at CPDs, we don't look at them in graphical form, but this is a typical shape. The high mark is around 23 percent on the black PDF line. That represents the probability of adding no claims as it tails out (the chart does go out to $1,000,000, but it's not shown). We convert back to the cumulative function, and that's the blue line. A Monte Carlo simulation would generate values, random numbers, from 0 to 1, and draw a line across to the point on the graph to the cumulative function, and go down. That would be the result for one life in one year. About 23 percent of the time, we'd expect a life to have no claims. There's also a miniscule possibility of hitting the $1,000,000 claim figure.

In Excel, it's easy to do. In Hendrickson Slide 2, page 5, on the left-hand side, I show the PDF and cumulative function. In Excel, you have to shift the cumulative function down for the look-ups to work correctly. This is how we're used to seeing PDFs, and this is the same PDF. It's just simplified so you can read the numbers. If the random number generated said a particular life was a 0.2, that would generate $0 of claim costs. If it were a 0.5, it would generate almost $7,000 of claims. For the 50-life group, you would do that 50 times. You'd add them up and divide by 50, and that's your average amount for that one iteration. You do that 10,000 times for each size of your group.

Here are the results (see Hendrickson Slide 1, page 6). You see the per member per year (PMPY) value at the top and the PMPM. Based on the central limit theorem, we can calculate not only the sample standard deviation, but a theoretical standard deviation because we can calculate a variance based on our underlying variance of our known PDF or assumed PDF. If we do that for each of these sizes of lives, you can see the difference. It's pretty close. Even for the 50- and 100-life group, where it is not normal, you can see that our standard deviations are good. You can see the minimum and maximum values, and as we go to the right, as you would expect, they tend to group closer to the expected value of $3,000.

I've also shown some margin amounts that if you took the 90th percentile result and related that to the average, suppose you're an HMO that wants to get out of business quickly and wants to charge groups based exactly on their risk margin based on this table. For a 50-life group, you would add a 62 percent margin on top of their expected claim costs. We wouldn't sell any business, so we don't do it that way. This gives you a background in how to look at the risk inherent in each of these sizes of groups.

If we graph the percentile of the results (see Hendrickson Slide 2, page 6), for 50 lives, they're spread out. We have more than $6,000 in annual claim costs for the 50-life group, and it goes down to where the variance is small. This is the same graph from earlier. For the 10,000-life group, it looks normal. The percentile rankings show that, as well.

If we go into a reinsurance situation, we get into all kinds of graphs. For the $50,000 specific reinsurance (see Hendrickson Slide 2, page 7), we have it set by group size. The red line for 50 lives looks nothing like a normal distribution. The way the reinsurance works is the annual claim cost values are going to be capped at certain values, so from the reinsurer's perspective, the $76,000 is going to become $76,000 – $50,000, or $26,000. Zero's above that, and you'd subtract the $50,000 off the other two, and that's going to be the distribution for the reinsurer. You can look at it from the company's perspective, as well, where you'd cap the values at $50,000, add the two together and get the same distribution. That's the background of how we're going to do this.

Where does it start to resume a more normal distribution? It starts to happen between 500 and 1,000, but you can make the argument 1,000 still doesn't look normal, but it's starting to take that general shape. By 2,500, it's starting to look normal, as by 5,000 and 10,000.

As you look at this, some numbers are going to jump off the screen at you and look like typos. They're not. If you look at the 90th percentile margin you'd need on specific reinsurance (see Hendrickson slide 1, page 8), for $100,000 of reinsurance on 50 lives, that would be a 12 percent margin. The reason goes back to the graph. This is for $50,000. If we were to run it at $100,000, what happens in the 50,000 lives means the only time there's ever going to be a claim paid is when it's more than 100,000 lives. If you remember the PDF, the probability of ever having a claim over $100,000 is small. For the $100,000 reinsurance level, the top mark on the 50-life group goes up to just short of 9,000 occurrences, so 9,000 of our 10,000 occurrences didn't have a specific claim. That's the reason these margin numbers don't make sense.

As you go farther out on the number of lives, they start to make sense. If we have any reinsurers in the audience, maybe they could let us know whether these are close. This was not designed to be a real-life sample. For the 90th and 95th percentiles, you can see what the margin levels would need to be at specific levels for reinsurance.

That was the reinsurer perspective. Now we can look at the ceding company's perspective. What happens to our claims if we had reinsurance. Again, I chose $50,000 of reinsurance to display the results (see Hendrickson slide 1, page 9). Everything shifts to the left from where it is without reinsurance, which you'd expect because the reinsurer's now paying part of the cost. It's squished things together and made them less variable, which you'd also expect, and made them more normal-looking. All the lines for the values with asterisks are the reinsured values. The solid values are the regular values.

On the previous slide, I showed you what the full claim-cost margin levels were for each of these sizes of lives. What happens if they were to purchase reinsurance from the ceding company's perspective? For the full claim costs, we go from a 62 percent margin to a 39 percent margin with $25,000 of reinsurance.