RECORD, Volume 30, No. 1*

Spring Meeting, Anaheim, CA
May 19–21, 2004

Session 21PD
Probability of LTC Ruin?

Track: Long-Term Care

Moderator: Janet L. Perrie

Panelists: Daniel Bret Cathcart
Amy Pahl
Janet L. Perrie

Summary: The NAIC rate stabilization model regulation requires LTC pricing actuaries to certify that rates are expected to be adequate under moderately adverse conditions. Currently most LTC actuaries test the sensitivity of rates on a deterministic basis. However, such exercises do not reveal the likeliness of a future rate increase. The panel demonstrates a technique that can help set confidence intervals regarding future rate increases. Specific topics include creating a stochastic model, deciding which variables are deterministic vs. stochastic, developing probability distributions for the stochastic variables, running the model and evaluating the reasonableness of results. Attendees gain a better understanding of the necessary steps to answer the question: How confident are you that rates will be adequate under moderately adverse conditions?

MS. JANET L. PERRIE: We hope to give you some ideas to avoid ruin by incorporating confidence intervals within your rate-setting. Our first speaker will be Dan Cathcart. Dan's with ERC and has been focused on LTC reinsurance for over 12 years. Next you'll hear from Amy Pahl. She's currently with Milliman USA and was previously at Life Care Assurance Company and Allianz Life, and she has been specifically in the LTC insurance industry for over seven years. I will conclude the presentation. I am with the Long Term Care Group. Prior to that, I was with Milliman USA. I have been working in the LTC industry for over seven years.

We hope this will be a very informative session for you. This was truly a collaborative effort with a lot of ideas and information shared among our three respective firms.

MR. DANIEL BRET CATHCART: As Janet mentioned, we want to show you a little more dynamic way to look at LTC pricing. Most of the industry has looked at it from a deterministic view. We want to move that along and show you a few methodologies to get to a stochastic model. Along the way, I'm going to talk a little bit about how actuaries are a lot like golfers.

When golfers walk up to a situation, there are a lot of conscious and subconscious assumptions that are taking place. They take a look at the direction the land is flowing, both on the green and in front of them. They might test the wind by throwing up some grass. They look to see if the flag or trees are moving around. Eventually they pick out the club, and it's the perfect club. It's going to put the ball in the hole with that perfect swing that they, of course, have mastered.

As pricing actuaries, we do much of the same thing. We have all these assumptions coming together, and, of course, if everything works out ideally, we'll come up with a target profitability that we can incorporate into our modeling. That's called deterministic modeling. We've figured out exactly what profitability we're targeting. Of course, a much more ideal methodology would involve being able to know, if we had 10,000 shots here, where we would end up. In this situation, what's the likelihood of being within 5 percent of our target? You'd rather know what the likelihood is of being 60 percent of the time on the green.

When you start talking about variations and distributions such as this, that's where stochastic modeling comes in. Eventually, of course, you have ruin. From a pricing actuary's standpoint, you'd probably set a hurdle, and that's where you have to raise rates. From a golfer's perspective, that's when you throw away the clubs.

It's not that easy to move from the deterministic model to a stochastic model. In fact, we're going to talk a lot during this demonstration, and we're actually only talking about one pricing cell. We are picking a cell that's most common—the age 62, comprehensive, $100 a day, 90-day elimination period, lifetime benefit period, with increases, and we're going to talk about a standard unisex rate.

Surrounding your pricing, similar to the deterministic model, you do have to come up with all the assumptions. You come up with specific assumptions for voluntary lapse, mortality and morbidity. We're going to break it out between incidence rates and severity. You also have to generate investment income scenarios, as well as incorporate expenses and profitability. We're going to walk through how you come up with your expecteds, but then how you expand that to come up with a variance. When you have expected and a variance, that's what generates your distribution assumption for each of these assumptions.

Starting with the voluntary lapse rate, this is an excerpt from the Life Insurance Marketing and Research Association (LIMRA)/SOA joint study that was done recently. We pulled out the section that shows the captive agency versus the brokerage agency, and for our testing we used the captive agency view as a baseline, but we did grade it out and make it smoother. Ultimately what we ended up using for the expected values by policy year was something like 4 percent first year, grading it down to 1 percent ultimate. There you have your deterministic view. You have your assumptions by policy year for lapse. But you have to turn that into a distribution for each of those policy years.

You have a 62-year-old. With each policy year, you have to have a different distribution. For lapse rates, you either lapse or you don't, and where that falls in is you really have a binomial distribution. The mean or the expected value for a binomial distribution is your probability of lapse times the number of simulations you have. I have used q just to try to stay consistent, since most people are used to using that for mortality. So, q is the probability of any of these assumptions that we talked about happening, such as incidence, mortality or lapse rates. But the theory is still the same.

In the end, for a binomial distribution, your simulations times the probability is the mean, and then to expand that for your variance, you have your simulations times pq or q times one minus q. That's your variance. So, taking all this information, again, by policy year, you have a different expected lapse rate incorporated in here. You have a different q by policy year for each one of the years for that 62-year-old. You take these assumptions—the mean, the variance and this binomial distribution—and you incorporate them into a stochastic model with this information, and that's where you build your model from. So that's one assumption.

Mortality is very similar. Again, you either die or you don't die, and in this situation I think you're more used to the q with mortality. But, again, we're using a binomial distribution. We went with the 1994 Group Annuity Mortality (GAM) Table, and, again, very simply you can put the information into a binomial distribution and create the mean and the variance. You have to do that for every single policy year and incorporate that into your model.

Next I'm going to talk a little bit about incidence rates. As I said, for morbidity we split up incidence rates versus severity. For incidence rates we had a fair amount of data for the eight to 10 policy year range, the select period, for the 62-year-old. We just used the straight experience results for the 62-year-old for those early durations, and Chart 1 shows those as marked by the x's. The x's are actual experience for each policy year.

Beyond the 10th year, since experience becomes limited beyond that, we have some experience, but we expanded on that by using all the issue ages but just dividing out the selection factors for those older issue ages. It was really more of an ultimate incidence rate that we used beyond the 10th year. Again, you can see where all the x's are. That's our experience. That's our best estimate based on experience of where incidents have fallen by policy year for a 62-year-old. But we wanted to smooth that out. So, using a regression analysis—and in this case it came out as an exponential function—we solved to create that smooth curve that you see there. For an exponential distribution it's e to the a times the policy year plus the b, and we solved for that a and b. Again, that formed this smooth curve that you see here. We did this both for male and female. I did talk about unisex rates earlier, but we did do all the modeling based on male and female and then combined the results in the end. Again, using this piece of the methodology, this was only driven to come out with smooth incidence rates by policy year.

In terms of the distribution itself, I'll go back to the same example as the prior two assumptions. Incidence rate either happens or it doesn't. So, you have the incidence rate by policy year, similarly your lapse and your mortality, and you can form that right into a binomial distribution again. Of course it's going to be different for every single policy year, but, again, you have your binomial distribution, and you have all the information you need to pile into your stochastic model.

FROM THE FLOOR: This question is going to show I haven't worked with this stuff in an awfully long time. For mortality, we can have a fairly high level of confidence. But lapses are all over the place, and we don't have a lot of experience after you get out a few years. Are you somehow going to catch that in your distribution here—that your variance for lapse will probably be much wider than your variance for mortality? I'm not sure how you're going to measure that or if you intend to measure that.

MR. CATHCART: As I said, it's all driven by the binomial distribution. So your variance is really driven by the rate itself. I think what you're getting at is there's a lot more experience on mortality than incidence rate or lapse rates. But the way that we did it, we assumed that our expected values are equivalently credible. The variance is driven by what you see here. It's the n times the p times the q. The lower the probability of any of the rates, whether it be mortality or lapse or incidence, it’s really the one-minus effect—or basically how low it is causes more variance to the result.

MS. AMY PAHL: I think the easy answer is that, no, it's not readily apparent that the variance that we see in real life for voluntary lapses is captured. Intuitively we have a mindset about that variance, but it's not explicitly captured, although we are convinced that the binomial distribution is correct to represent the voluntary lapse assumption. The point you make I don't think is directly addressed, so it would be a limitation. I think maybe what you would do is have to somehow link that in with correlations between assumptions.

FROM THE FLOOR: You reached a certain year. You're in year seven, let's say, and without regard to anything that has happened before, except the fact that the person hasn't died, hasn't lapsed, and you have a mortality rate that you think is right. Well, the mean is the mortality rate.

MR. CATHCART: The mean is the mortality rate for that policy year.

FROM THE FLOOR: Likewise for lapses.

MR. CATHCART: Right, and likewise the lapses. It's a good point that you're bringing up. If this is our best estimate in terms of moving toward a model like this, we are going with the assumption that binomial is a decent fit.

FROM THE FLOOR: You don't really talk about standard deviation then when you're talking about a binomial distribution?

MR. CATHCART: The standard deviation is the square root of the variance. So, it's all driven by the binomial distribution in the formulas there.

FROM THE FLOOR: You're saying that your expected value is right. If we look at mortality, we have a good idea of what mortality is for 1994. What we don't know is what mortality is going to be in 2014. So, you were trying to do something that looked at what would happen with various types of mortality improvement. You would have to change your expected values for this.

MR. CATHCART: Definitely.

FROM THE FLOOR: An extension of this might be a two-stage simulation where you first simulate the expected value or the q, and that could reflect observed volatility in lapse rates that's greater than observed volatility in incidence rates. Then after you've simulated the future lapse or incidence rates, you simulate the actual experience without that value, taking into account the sampling error that you're getting out of this.

MR. CATHCART: Right.

FROM THE FLOOR: The simulation is kind of perceived from the fact that you know what the mean is. If you have a bad number in there, you're just going to have bad outcomes because you're proceeding with that. There isn't any way, simulation-wise, to correct that information.

MR. CATHCART: You're right. You are basing this all on decent input that you're putting into the model.

FROM THE FLOOR: The '94 table is not the correct mortality underlying this population. Then you just have bad answers.

MR. CATHCART: It is critical that you get your data and your inputs right. Let's say my '94 mortality isn't perfect. At least in the end it give you some idea of the variance and where LTC can go, because I haven't seen much that shows that. What's the probability, if you price it to a 15 percent return on equity (ROE), that you come out losing money? You'll get some sense with what we've done, but you're right. As you fine-tune it, include covariances, you will get a different answer, but this is at least a start to give an idea of how the distribution looks from an LTC pricing perspective.

FROM THE FLOOR: Variable n is sort of a random variable itself.

MR. CATHCART: Well, n is the number of simulations. I'll leave that until later.

FROM THE FLOOR: Variable n starts as the initial amount issued but after five years n has been decreased by all the other decrements.

MR. CATHCART: You'd want to test how much your results vary by how much and what your simulations are. We'll talk a little bit more about that later.

FROM THE FLOOR: It seems to me that n should be the number of policies.

MS. PAHL: It's the number of trials. Each duration you're doing 10,000 simulations. It's not like you issue 10,000, and it runs off. You have 10,000 in every duration. You don't have the effect of decrements on that n.

MR. CATHCART: Yes, it's 10,000 simulations for each distribution for each policy year. If you run 10,000, in the end you'll see a distribution of where that one policy could fall out. It's going to show you have a 5 percent chance of being down here or 5 percent chance of here. The idea of doing many simulations is to build this probability distribution, so then you know with one policy where you're likely to fall. That's the idea of it.

MS. PERRIE: But when you say one policy, it's not more like a Bernoulli where you have one person that starts, and that person either lapses or dies or nothing. It's not that kind of one person. It's still just the general one person who has the 2 percent chance of lapsing.

MS. PAHL: You can think of it exactly the same way you're doing your deterministic pricing. You issue a block of policies or one policy. You ratio it to $1 million of premium, or whatever you want to call the starting amount. You're making certain assumptions, and you're projecting it. We're doing the same thing except now for each policy year value which had been on a deterministic basis, we're now saying we're going to run 10,000 simulations and draw from a distribution. So, instead of having 4.25 percent first year every single time deterministically, we're going to have that as the expected value, but then a binomial distribution from which we'll pick. You'll end up with 10,000 different values for first-year lapse rates. In the second year, you'll run 10,000 simulations, you'll have your expected value being the deterministic lapse rate and you'll pick from that distribution. On average it should come back to the mean, but you're going to get 10,000 different second-year lapse rates. Does that help clarify what the extension from the deterministic is?

FROM THE FLOOR: Another way of doing it is you issue a person, you simulate, and you say, oh, this one died in year seven, then you do another simulation, and, oh, this one went on claim in year nine, and come up with a distribution that way.

MS. PERRIE: That's more of the Bernoulli.

FROM THE FLOOR: Yes, but that's not what you did. You came up with a distribution for each decrement.

MS. PERRIE: Right.

MR. CATHCART: Yes.

MS. PAHL: Right.

FROM THE FLOOR: What they're saying is you're not recommending that anybody do this; this is just to illustrate the concepts.

MS. PAHL: I wouldn't necessarily say that it's not something that we believe is correct or an appropriate way to go about it. It's maybe only one of a few different ways, but I don't know that it would be something we wouldn't recommend.

MR. CATHCART: I would recommend at least taking a look at this way and fine-tuning it. As Amy said, there are other ways to do it which have been discussed, but basically we're pointing out that just looking deterministically doesn't give you a lot of insight, and you can learn a lot more by looking at more of a distribution view, a variance view.

MS. PAHL: It's also important to keep in mind that there are a lot of constraints on this process. The first is whether we even have software tools to do this for us. Even though we can dream theoretically how all of this is ideally supposed to work, there are the practical limitations. What we're trying to demonstrate here is something that you can do within the constraints we currently have to live with, and until there are better systems, better tools or better data available to us, we believe that this is an approach that gives you some more robust results without completely disregarding some of the simplifying assumptions. There are some things that we admittedly have to assume. Those are limitations. One of them that I mentioned and that will be mentioned again is that there's no correlation between assumptions. Clearly there is a correlation between assumptions, but we don't have a tool to incorporate that correlation coefficient at this point. The systems aren't sophisticated enough to do it. I wish we had all the answers.

FROM THE FLOOR: What you're saying is this isn't the ideal, but, based upon current constraints, it's a way to get a more insightful look at what is going on.

MR. CATHCART: Correct.

MS. PAHL: Exactly. I think you'll find by the time we finish our remarks that it is quite a bit more insightful and more revealing as far as what the risk profile is relative to what we learn when we price deterministically.

Up to this point, everything has been a binomial distribution. But one of the actual complications that we had to deal with going in is that for the most part people use claim costs, and putting a distribution on claim costs is not really correct because the two different pieces are so distinctively different. So we did split out incidence and severity, I believe appropriately—incidence being a binomial distribution, fairly straightforward from a stochastic standpoint.

Severity was a lot different to deal with. We obviously wanted to base it on company experience, and we had two batches of data. We looked at both of those. The one we ended up using and reflecting in the analysis was a very small block of claims, as it turns out. There were about 768 claims. About half of them were closed. About half of them were open. For those open claims, obviously we had to project or estimate the remaining expected length of stay using expected continuance. You're still doing that on a deterministic basis. For purposes of both incidence and severity, we defined the length of stay or the incidence of claim being that surviving the elimination period. So, this is not that you're going to actually be activities-of-daily-living (ADL) deficient or cognitively impaired, but you're going to actually qualify for benefits and be on claim at the end of your elimination period. The length of stay that we're talking about here is the first payment through the last payment.

We did, in the case of both incidence and severity, take out all the claims that did not match our one-cell profile. We only have 90-day elimination period, lifetime. We did include both compound cost-of-living adjustment (COLA) and no COLA because we don't think that affects severity. I don't think we assumed it would affect incidence, either. We included both in development of the assumptions for incidence and severity.

We thought the more straightforward approach would just take these data points and put them to a known distribution. See if we could name it. The name of it then develops into a mean and a standard deviation generally. We tried to do that with both sets of data. Incidentally, the second set of data had a few more claims. It was a group block. It was younger attained ages, generally. The results were not consistent with this block, which was an individual, more traditional, 60-some-year-old average issue age. So instead of combining the data, we just went with the one set that was more representative of the 62-year-old that we're trying to demonstrate. But it was interesting because we looked at both sets, tried to fit both sets to a known distribution and failed miserably in both instances. The fact is it just doesn't look or act like a distribution that's common and built into the systems. Incidentally, we did use a pricing software system that does have stochastic capabilities, but none of those distributions worked that were in the pricing system.

We ended up just plotting the data and developing our own histogram. Fortunately, the system could take that and draw from that distribution, even though it wasn't one that was well-defined. Chart 2 shows the histogram. It's interesting having the high early frequency of short claims—frequency meaning the number of claims that had a length in that interval. The first tall blue line is the number of claims between zero and 100 days. Then each of these lines represents another increment of 100 days. The hump centers around about 1600 days, and the mean of this turned out to be about 1171 days.

FROM THE FLOOR: You have two groups of people, at least two—nursing homes, one of which includes the acute care and one of which includes the LTC, and you have a benefit that's basically getting paid for the acute care as well as LTC.

MS. PAHL: Yes, I think you're right.

MR. CATHCART: We agree, and we thought this looked exactly like what we expected. This is something that seems very sensible to us. The question was: How do you model it?

MS. PAHL: Right.

FROM THE FLOOR: Yes, but it's surprising that there aren't more just deaths alone from the conditions that lead to confinement that would make it look more regular. What is the scale of the bottom axis?

MS. PAHL: Each line is 100 days. So, for example, the last, tiny, little line at the end here is 4500 days.

FROM THE FLOOR: Between 4401 and 4500.

MS. PAHL: Exactly. Like I said, the high mark there, which is kind of right in the middle of that hump, is 1600—1501 to 1600 days.

FROM THE FLOOR: This is all after the 90-day elimination period.

MS. PAHL: Yes.

FROM THE FLOOR: So, this is 91 through 191.

MS. PAHL: Yes. You'd have to add the 90-day elimination period onto this length to get the total length of qualification on claim. This is comprehensive, so it's all care on a comprehensive policy. We did throw out the ones that were single day, obviously non-forfeiture or what have you, but anything that was greater than a day, so it could have been even a few days in adult day care. You're going to get some of that kind of stuff in here as well, which would explain that early spike or the early few tall lines.

FROM THE FLOOR: Did you try doing this separately for home care versus facility?

MS. PAHL: We did not.

FROM THE FLOOR: Is this all attained ages here in excess of 62?

MS. PAHL: That's a good question. I can't remember if I took out under-62 or if I just said that for this purpose let's just use all the claims we have because we don't have very many. I think the answer is that I didn't throw out any of the young ones. But this block was issued over a 14-year span. It was your traditional, individual, 60-something issue age block of business. You're going to get some young ones but not a disproportionate number.

FROM THE FLOOR: For the claims that are still open, did you plot them as having a length of stay equal to the duration as of the latest time you had observed them, so they'd been open for two years?

MS. PAHL: No. We said that there are two years, but there's also an expected continued portion of that claim. We used the experience continuance table to anticipate the remaining duration. You're deterministically developing how much of that open claim will be in the future, and then the total claim, both the actual past paid portion and the expected remaining, is plotted here together. This graph represents all 768 claims.

We did impose some adjustments and constraints on this. We didn't just take 1171 days as the mean length of stay and apply it across the board. We did recognize that as people aged, their length of stay was declining. We also recognized the fact that there's some constraint or there's some limitation on your length of stay just because of mortality. We didn't want to have the 95-year-old claimant lasting 1171 days; it's not a realistic assumption. We did grade off the length of stay 2 percent per year by claim and age, so it was reducing over time. Then also we layered on another constraint that the total claim length of stay would be equal to or less than the expected lifetime under four times the U.S. life mortality table, which happens to be one we've experienced that has some relationship to the disabled life mortality. Basically, you identify what you think your disabled life mortality is, and you make sure you use that as another length-of-stay constraint because the claimants are going to die at a certain rate once they're disabled, and you don't want to have them staying on because you have not imposed that death rate. Those were two additional adjustments.

Fortunately, the system that we used could take this histogram and draw from it for each year again as it developed. It pulled from the binomial distribution and incidence rate, and it also pulled from this distribution a length of stay. For each policy year projection point in your 40- or 50-year projection, you're doing that 10,000 times. You now have a histogram of claim costs, if you will, which we didn't look at or plot or anything (although it'd be kind of interesting to see what you get when you do that), but that's theoretically how it's working. We did a lot of this together. The fact that I'm presenting this doesn't mean I did all of this. It's just that the data that we used ended up being data that I had been working with.

FROM THE FLOOR: You have the incidence with the binomial distribution, and you have the severity. You're generating a random number and picking how much incidence you had and how much severity you had and doing that 10,000 times, which takes at least 20,000 random numbers—well, more than that, but you need one set for the incidence. You need a totally separate set for the severity.

MR. CATHCART: The model did that itself.

MS. PAHL: The projection software had that capability.

FROM THE FLOOR: Does your length-of-stay data include termination for death?

MS. PAHL: Yes. The length-of-stay data reflects all reasons for leaving claim status. We didn't differentiate whether it was a recovery or a death. They were there on claim for a certain length of time, then they weren't.

FROM THE FLOOR: Your average length of stay includes terminations for death, so it has older people who would possibly have a shorter claim.

MS. PAHL: Yes. That's why we impose the second constraint of making sure that the length of stay was not more than the expected lifetime of a certain disabled life mortality table.

We took all the data, we came up with the average and then we said that that works on average, but as people age there are practical constraints, and there's not the expectation that people are going to stay on claim on average as they get older.

MR. CATHCART: I have just an example of that. It wouldn't really come through for a 62-year-old, but, say, for the incidence rate—it came through on age 99. That's where we restricted it. For the 99-year-old plugged into this severity, we limited it so that it couldn't pick from this chart way at the far right.

FROM THE FLOOR: But the chart already reflected such claims. In other words, you're truncating your older ones—the average of all of your ages now is no longer going to be that.

MR. CATHCART: Slightly.

MS. PAHL: The other thing is that there's no actual experience in here for claims of that attained age.

MR. CATHCART: Right. This is based on the earlier age claims.

MS. PAHL: The few claims that would be in what we'd consider those higher attained ages that you'd have to constrain would be so minimal relative to the few that we had anyway that we did recognize that putting that constraint on there would remove it from coming back to this average, yet you couldn't get a deterministic result that made any sense if you didn't put constraints on it.

FROM THE FLOOR: But is the constraint just slicing off beyond a certain duration? In other words, are you getting claims at age 93? I'm going to keep the bars exactly the way they are until a certain duration and then slice it off. Did you actually attempt that?

MS. PERRIE: We did this on a 62-year-old, and so the first policy year we used this chart straight out with no adjustments. In that way, truncating it using the disabled life table didn't do anything. We didn't do any adjustments. Then the next year, when the person is 63, we took the mean of 1171, reduced it 2 percent because we picked a rate to decline the length of stay over time because the 102-year-old isn't going to have the same average length of stay as the 62-year-old. Then we imposed the life limit. The life limit didn't hit at that age. The expected lifetime limit didn't hit for 25 years or something, and so that didn't come into play until much later, at which point we don't even have as many lives coming through and so on. It wasn't a major thing to do, but we thought it added validity to the model. Remember we're using the '94 table where certain death isn't until 120.

FROM THE FLOOR: So the mean is decreased by 2 percent compound for each year.

MS. PERRIE: For attained age over age 62.

FROM THE FLOOR: Once you get that mean, do you have to construct a new histogram?

MS. PERRIE: For the deterministic model, the severity was the 1171 with all those adjustments. To go to the stochastic, we used that shape, randomly picked our number and then adjusted that new number for the new mean at whatever attained age we were. So, for the 72-year-old, let's say the overall adjustment was a 0.9. I would pluck off whatever number that was and then multiply it by 0.9, essentially to squish it down to recognize that our mean has decreased at that age. We're just using a cumulative distribution function to randomly pick the number.

FROM THE FLOOR: Basically you kept the same look.

MR. CATHCART: Yes.

MS. PERRIE: Yes.

MS. PAHL: We kept the same look proportionately. As we would do on a deterministic basis, the stochastic formula would be to pluck the incidence rate from the distribution and do the same for severity. Then we had an assumption for severity intensity that was not a stochastic assumption. It was just a standard constant that we set as this formula describes. In doing so, we made some assumptions about the proportion of claims within that comprehensive coverage—what would be facility and what portion would be home care and then how many days a week they would be receiving services and, therefore, claim dollars, and what the dollar or the 90 cents refers to the portion of the maximum that they would be utilizing. We're saying that somebody who's in a facility is going to need it every day and is going to spend the maximum. Somebody who's at home is only going to use it four-sevenths of the time, and they're only going to spend 90 cents on the dollar, which did not wear off but we thought was reasonable since it's compound inflation protection, and that margin is likely going to stay at least for a good portion of the projection period, if not the entire projection period, depending upon your assumed inflation rate on the cost of services. We combined these three parts into a claim cost for each of the projection years. This shouldn't look too unfamiliar.

FROM THE FLOOR: What is the four-sevenths for?

MS. PAHL: That reflects that four days out of seven you'd receive home care.

FROM THE FLOOR: Of course.

MS. PAHL: We did make investment income stochastic as well. For this one we relied on canned software, and if you've used actuarial projection software very much for life and annuities, they've been doing this for years. There's generally an interest rate stochastic generator on almost any actuarial projection software program, and that's what we used. The way that works is it takes a random number—a lognormal distribution or lognormal process with a reversion to a mean—and you get interest rates for either monthly or annually, however you want it, based on an initial yield curve, which we assumed was year-end. Because obviously LTC insurers don't invest in risk-free assets, at least not entirely, we did add a spread over the risk-free rate to represent what a common LTC block of assets would be.

We did target the 5.5 percent rate, which is what we chose to use on a deterministic basis. The stochastic process needed to have an expected value of 5.5 percent as well, so we have a mix of business that did that. The way the interest rate generator worked is it takes every point on the yield curve, or every point on the yield curve you want to give it, and will stochastically generate 40 years monthly or annually. We, of course, didn't want to use every point on the yield curve. We chose three different points in proportion so that, again, it was intended to represent what an LTC insurer might really use for assets to back their LTC. We used a combination of the one-year maturity to seven-year maturity and the 20-year rate. Has anybody done a lot of this on a non-LTC product? Does all that sound familiar from what you've done on life or annuities? There isn't intended to be anything special or different about this.

Expense is one that I don't think is random. It is set deterministically for our analysis here, and it was just designed to represent what a current product being developed today might look like. Whether it is or isn't is for everybody to determine themselves, but I think it does that job of representing what common expense structure might be.

FROM THE FLOOR: If I were a smart aleck, I would say that your inflation would have been linked to your investment in some way.

MS. PAHL: That's where I say our limitation is that we don't have correlation between the assumptions, and that is one of the recognized drawbacks of the analysis.

Then, as would be done with a regular pricing exercise, you'd pick a profit target. To keep it realistic yet simple (we didn't want to use anything that was GAAP-related), we used statutory internal rate of return (IRR) to determine the premium on a deterministic basis. That's probably consistent with what most people are doing. Whether it's 15 percent or something different certainly could be a variable, but it seems realistic. Then, to keep it a little simpler on the stochastic side, we converted that IRR to a present value of distributable earnings as a percentage of present value premium. It's an easier profit measure to calculate when you're doing it 10,000 times, as well as the fact that IRR can be nonsensical if you have all these sign changes which could in theory happen under a stochastic process. For use in the stochastic analysis, on a deterministic basis we equated different IRRs to a present value of distributable earnings. This was our assumption: 15 percent translates into a 9.7 percent present value of distributable earnings and so on.

To summarize, certain assumptions—the lapse rates, the mortality and the incidence—are binominal. It's like flipping a coin. It's either heads or tails, lapse or not lapse, die or not die, go on claim or not go on claim. Those are pretty straightforward. The severity was from the histogram, investment income from a lognormal distribution and expenses set deterministically. Obviously there are assumptions here that we didn't mention. The reserves, of course, are going to be deterministic. Those aren't going to change even under the 10,000 different scenarios because you wouldn't actually be changing reserves under different scenarios. All the others are deterministic. We used the common, everyday risk-based capital (RBC), target surplus and tax assumptions. The intent was to come up with a premium that would be consistent with what people are developing today for a 62-year-old for this benefit structure. As I mentioned before, at least one of the limitations in the analysis is that there's no correlation between stochastic variables. Now we'll present the results.

MS. PERRIE: Using all those assumptions that Amy and Dan outlined, I input them into our pricing model and targeting that 15 percent IRR, I calculated a premium of $3734. It might immediately look a little high, but it is a standard rate cell with comprehensive benefits and conservative assumptions. Chart 3 outlines the components of premium. Discounting the cash-flow streams at 5.5 percent, which again is our pricing investment earnings rate, results in this allocation of the premium. Again, this is just the single age-62 rate cell. This might not look exactly like a profit study result where you're aggregating many, many cells, but hopefully it's not too far off from what you're accustomed to seeing.

At this point this covers best estimate assumptions. But, of course, under the rate stability regulation adopted by the NAIC, the next step is to consider moderately adverse conditions (MAC). According to the model regulation, the actuary must certify that "the premium rate schedule is sufficient to cover anticipated costs under moderately adverse experience, and the premium rate schedule is reasonably expected to be sustainable over the life of the form with no future premium increases anticipated." The exact certification that we all have to sign off on is straight and clear, but, of course, what's not clear or defined anywhere is "moderately adverse experience" or "reasonably expected." We are not going to propose any steadfast definitions for that either. We just hope to offer something to consider as you're performing this analysis on your own.

Following deterministic methods that we might find common in the marketplace, we subjected the $3734 premium to MAC. Specifically we looked at lapse, investment earnings and morbidity. What we're saying with these tests is: If this deviation occurs, will we refrain from seeking a rate increase? As long as we're satisfied with the resulting profit, we're comfortable certifying that the rates are adequate under MAC. For lapse, we considered voluntary lapse at 80 percent of pricing expectations to be a MAC.

FROM THE FLOOR: These are three different tests.

MS. PERRIE: These are independent tests, one at a time.

FROM THE FLOOR: So you vary lapse and leave everything else alone.

MS. PERRIE: Yes.

FROM THE FLOOR: But this is deterministic.

MS. PERRIE: Yes. For lapse, we looked at 80 percent of pricing expectations. That only dropped our IRR from 15 percent to 14 percent, but, if you remember from earlier, we're starting at 4 percent and grading down to an ultimate lapse rate of 1 percent, and so dropping the 20 percent didn't impact our plan very much at all. But plans that are only a few years older than now that have higher ultimate lapse rates would have a stronger impact from this assumption. From our perspective, that dropping lapse is 20 percent and it only dropped our IRR one point. No, we wouldn't consider a rate increase on that.

Resetting the lapses back to 100 percent, we looked at investment earnings being 50 basis points lower than our pricing expectation. Again, that only dropped our IRR about two points, and we wouldn't consider a rate increase under that condition either.

Finally we looked at whether morbidity was 10 percent higher than expected. Again, at this point morbidity is a combination of incidence, severity and service intensity. If those three components together result in higher morbidity of 10 percent, would we expect a rate increase? Since that only dropped our IRR about three points, no, we don't consider that severe enough to have a rate increase. So, the way we're looking at things these days, if this were the extent of the universe of the conditions you are going to test for moderately adverse, yes, I feel comfortable that this is adequate for moderately adverse and I would be comfortable signing the actuarial certification.

Now, I imagine this process (maybe not these exact conditions) sounds familiar or at least reasonable to many of you, but what is this really telling us? All it's saying is that if any one of these deviations occurs in isolation, is our premium okay? But we all know that our best estimate assumptions are rarely realized, and not only that, but even if we are right somewhere, we're not off on only one variable. With this in mind, we decided it was better to consider the likelihood of achieving our profit target and, maybe more importantly, not the single cause of what might trigger a rate increase, but the probability that we might require a rate increase.

I'm going on to the stochastic analysis. We first tried 1000 scenarios, and, again, that's varying the lapse, the mortality, the incidence, the length of stay and investment earnings over the 1000 scenarios. Now, unfortunately, this histogram was a bit jagged and not very bell-shaped. We would have liked a nicer bell, so we increased our trials to 10,000. On an aggregate basis, the mean of each component of premium very nicely comes back to our deterministic starting point.

However, looking at the resulting histogram (Chart 4), it shows a much different picture than just looking at the prior table. There's severe potential variability that can occur. Even after running the 10,000 scenarios, we don't really have a smooth curve. There are still those few drops at about 6 and 18 percent. But, again, I didn't run it multiple times to find out if this was reproducible or to isolate the cause. Again, we're not considering correlation or covariance among the stochastic variables. With those limitations recognized, we decided to move on. But 10,000 trials were a little time-consuming, and we didn't deem it worthwhile to increase it beyond that.

FROM THE FLOOR: You said earlier that you did incidence, I think, for males and females. Deterministically you would have had to make an assumption in terms of how many male and how many female. What do you do stochastically for this?

MS. PERRIE: That just stayed 50/50. We didn't vary that.

Using the previous graph as a measure of possible returns, we now consider the probability of a rate increase being triggered. We chose estimated lifetime profit versus expected to be our trigger, and we considered two minimum thresholds—one, that a company might decide that they really want to have double-digit IRR returns, and the other being that they would be satisfied if they earned half of what they originally priced for. Again, the IRRs were translated to percentage of premium for ease.

If a double-digit IRR is going to be considered to be the minimum threshold, the 10 percent IRR translates to 5.5 percent of present value of distributable earnings as a percentage of collected premium. Using the $3734 premium and the 10,000 scenarios that we previously ran, this results in about a 30 percent chance that our profit will fall below 5.5 percent or, said another way, there's a 30 percent chance that we would need a rate increase. Similarly, if half of the IRR from the original pricing assumption is used as the minimum threshold, the 7.5 percent IRR translates to 2.8 percent of premium, and, again, those translations were back from the deterministic model way at the beginning. In this case the probability of needing a rate increase falls to about 18 percent.

FROM THE FLOOR: But this is all saying that your assumptions are correct. You're just experiencing random deviations.

MR. CATHCART: Sure.

FROM THE FLOOR: You would really have a larger probability of a rate increase because your assumptions might actually be wrong.

MS. PERRIE: You have to start somewhere.

FROM THE FLOOR: In other words, if your assumptions were actually wrong, you would also see the probability of the rate increase become much higher.

MS. PERRIE: Well, if the mean of your assumptions was entirely off.

FROM THE FLOOR: Yes. This is saying in some sense that assumptions are being realized.

MS. PERRIE: Right. But overall we're hoping our assumptions come back to the mean.

MR. CATHCART: But remember we're trying to model stochastic versus deterministic. You have the same issue with your deterministic modeling where you're using a deterministic average as the basis to move onto the stochastic model.

MS. PERRIE: At this point we ask ourselves: Is this good enough? Will management be happy that we have a 30 percent or an 18 percent chance of needing a rate increase?

FROM THE FLOOR: Within that 30 percent there would be the situation where you need a 1 percent rate increase.

MS. PERRIE: Oh, sure, depending on what you're trying to bring your rate increase back up to.

FROM THE FLOOR: Yes. That stochastic set would say I didn't meet my requirement. I need a 1 percent increase to meet my 10 percent IRR.

MS. PERRIE: Yes.

MS. PAHL: But if you were going to go after a rate increase, you might not go back to get just 10 percent. You might go back to get your original profit.

FROM THE FLOOR: It would seem to me one test you would want to say here is that I want essentially a 10 percent IRR, but also how many of these would be at least a 10 percent rate increase?

MS. PAHL: Right. We don't know the answer to that with this data.

MR. CATHCART: You're setting your threshold, probably that double-digit threshold. You're probably saying if it goes below that double-digit threshold, then you'll incorporate a 10 percent increase. We're not necessarily trying to get back to the 10 percent. Somewhere you're setting your threshold of when it is worth going forward with the rate increase. If you're going to do a rate increase, you're going to do at least 10 percent.

FROM THE FLOOR: This is fascinating, Janet. One question that I just thought of now and I should have asked before is, you're running this 10,000 times, but this is really one cell, correct?

MS. PERRIE: Right.

FROM THE FLOOR: In reality, it seems to me that what you have to do is look at a sizable block of business. On some of these items there's reversion to the mean. If your assumption is right, one person's variance will be out here, but if you have 10,000 or 100,000 people, you have somebody else who isn't making the claim, and at least where morbidity is concerned they'll revert to the mean. Where mortality is concerned they'll revert to the mean. Now, with interest that's not true. Interest is going one way or the other on your whole group. But your bell curve will be narrower if you have a large group on some of the variables. I know on mortality and morbidity that would be true. For interest it's not true. For lapse, I think it's not true.

FROM THE FLOOR: How did you take the 10,000 trials and model them within assumptions so that you had one offsetting the other? Or did you run it just for an aggregate set of assumptions for each trial?

MS. PERRIE: Just the aggregate set, yes.

FROM THE FLOOR: So you'd have a much narrower expectation if you were independently varying each assumption so that some offset each other.

MS. PERRIE: They were all varying at the same time. Some were offsetting.

FROM THE FLOOR: I thought the trials covered all the durations.

MS. PERRIE: Yes.

FROM THE FLOOR: Is that correct? But if it covers all the durations, and you get to the end, isn't it too late to ask for a rate increase?

MS. PERRIE: I guess at that point you say that your experience has been this way, and you're assuming it's going to continue that way.

FROM THE FLOOR: That's an easy answer. The probability of measuring is the probability that you asked for a rate increase.

MS. PERRIE: Will management be happy with this? Assuming the answer is no, they would really prefer that that chance of a rate increase be lower than 30 or 18, we next have to determine, how much of a premium increase do we need in order to make management happy, but still have a marketable premium? You could increase the premium so much that there'd be very little chance that you would lose and need a rate increase, but there's practicality.

FROM THE FLOOR: If you look at life insurance experience where they have a full cash value, long-run lapse rates are something on the order of 1 percent. Why should LTC have higher lapse rates than life insurance?

MS. PERRIE: We're not modeling that. We're modeling that as a 1 percent, too. Are you saying that it should be even lower than a one ultimate?

FROM THE FLOOR: I would predict that eventually it will be less than 1 percent.

MS. PERRIE: Next is the premium adjustment. Again, we start with our baseline 15 percent threshold that we priced with that came back with a $3734 premium, and so we'll meet that half the time. Now, the $3734 premium satisfied the 10 percent IRR 70 percent of the time. We looked at what we needed to do to the premium to increase that to an 80 or a 90 percent probability that we would not need a rate increase. Similarly, the 7.5 percent IRR minimum threshold was met with a $3734 premium 80 percent of the time. We wanted to look at what premium would bring that probability up to 90 or 95 percent. As it turned out, the premium of $3950 satisfied the first step of each IRR threshold. That translates over to a 5.8 percent premium increase. However, to improve to the next step up to the 10 percent IRR 90 percent of the time or the 7.5 percent 95 percent of the time required a 12.5 percent increase. That is the increase required now before ever issuing it.

Each company will look at this table differently. One company might look at this and say that it's still happy with the $3734, and it can live with that. Another company might look at it and say that it won't possibly go out without that $4200 premium because it really doesn't want to ever be contemplating a rate increase, or somewhere in between. Each company’s situation will be different, and they'll have different starting assumptions and different profit targets and different thresholds. This is just one possible scenario that is not perfect but was the best we could come up with right now. We hope we gave you something to think about and some worthwhile suggestions for incorporating stochastic analysis in your MAC testing going forward.

FROM THE FLOOR: It seems to me if you present these results to company management, you have to start the discussion with the caveat that this is based on historical utilization, and the probabilities that you're presenting assume the future looks like the past. A natural question that you might get from management would be, how much more variability should I expect from the results given that the future may not look like the past?

MS. PERRIE: If you really wanted to get fancy, you could incorporate some of that future change, improvement, whatever, in your stochastic analysis in order to get a better judgment value. You're not going to go out and price the product explicitly having a lot of change, but you could do a two-step analysis because, you're right, the future won't be the same as the past, and you can only learn from where you've been.

MR. CATHCART: Yes. Remember, you're controlling every single policy year. You can price the future as well as you can in the deterministic setting. You can price whatever you foresee changing from the experience into the future with the same methodology.

FROM THE FLOOR: Let's say you're an LTC-only company. This would also be a good suggestion for the actuarial opinion in the sense of determining the probability of ruin not just in terms of profits but how much surplus you need basically to reduce your probability of going insolvent down to 5 percent or 0 percent or something like that. I assume you could use this stuff to draw those kinds of conclusions also.

MS. PERRIE: Oh, sure.

FROM THE FLOOR: First of all, I'd like to suggest that this is more of an exercise in the probability of needing to do a rate increase rather than the probability of ruin because when we get into a conversation about surplus and so forth, what we're really talking about is that in a noncancelable environment this might well have some predictive ability with what levels of surplus you have to support it, but in an area where if you exceed moderately adverse you still can go after those rate increases, then, of course, you're also able to avoid any ruin scenarios that may apply, and that there certainly would be a heavy management input to this process. The other thing I was going to say is that I suspect you can make very subtle changes in your mean on morbidity, in particular, and certainly on all the other assumptions, and wind up with very different probabilities of rate increases so that, in fact, you could manage that if that were your goal. The final thing I was going to say is that when you put those higher premiums on, and you said, "Now I only have a 10 percent chance of a rate increase," it was really equivalent to saying, "I really want an expected ROE of 18 instead of 15 or some other higher number." Where you set that ROE factor to start off with and how far down that range you're willing to go is going to have total, what I'll call overriding, effect on whether or not you ever would have an anticipated rate increase.

MS. PERRIE: Yes.

FROM THE FLOOR: Fortunately none of my probability professors are still in the room here, but if you did this for 10 quinquennial ages or whatever here, you really wouldn't change the shape of that curve. These things are all independent, but you'd think that the curve would get a little narrower.

MR. CATHCART: I think it would be narrower for the older ages. I think it would be wider for the younger ages. There is just one thing I want to add because I actually did show this to senior management to give them an idea of what LTC looks like, and we're all pretty involved with LTC. I come from ERC, which is a property and casualty (P&C) company. This was a great message when they saw this. P&C obviously is all over the place. From a stability standpoint, this actually gave a lot more credibility to LTC when I presented it, showing it that, yes, I'm pricing to a 15, which everybody can say they're pricing to a 15, but this variance, even though there are a lot of caveats to it, there's a lot of fine-tuning. It gave a much better picture of where results could come out. I would just suggest that there's a lot of power to this for those that aren't involved in LTC.

FROM THE FLOOR: Did you happen to look at the tails in particular, trying to understand if there was a certain cause or a certain assumption that was driving the high variances in the tails? If lapses aren't a big deal, when you see a lot of low-lapse or high-lapse scenarios and with the interest rates, you see a pattern in the tails when you look at those scenarios. Did you go backwards and look at these and try to categorize which scenarios were driving which kinds of rates of return?

MR. CATHCART: Not in the time frame, but that's a great idea.

MS. PERRIE: It would be an interesting thing to do, to just do one of them at a time or one building on the other. It took an hour to run the 10,000.

FROM THE FLOOR: On the right-hand tail, which is where everything is going very nicely, you would also run into your not paying out according to the loss ratio.

MS. PERRIE: We're beyond the loss ratio at this point. We're starting with rate stability, and the loss ratios aren't a factor.

FROM THE FLOOR: For small substandard blocks of business, how do you determine a rate increase that would make sure that it provides for adequacy for MAC?

MS. PAHL: That's a tough one to answer. I don't think you answer it with this kind of analysis. If it's performing poorly, and it's substandard, that's the classic spiral situation.

FROM THE FLOOR: In spite of all the questions, I think this is a major step forward from what I've seen. Working on the Practice Note that came out, we said that we would like to do more stochastic and have some range for probability distribution, but we really didn't see anything. I want to congratulate you on moving forward. I think this is a good first step. I think the questions that we asked are because we're interested, not that we think this isn't something useful.

Chart 1

Incidence Regression

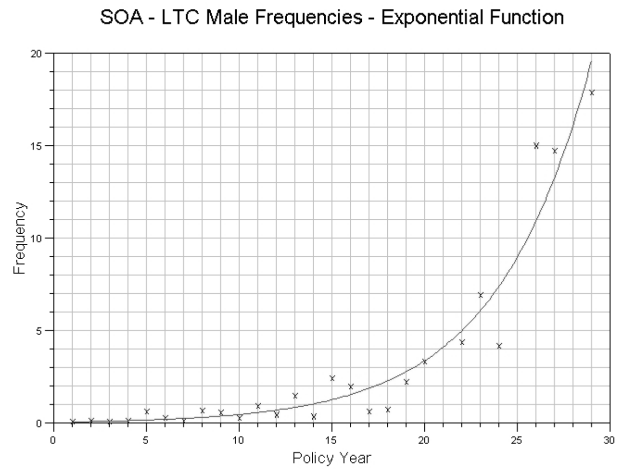

$e^{\mathbf{a} \times \text{policy year} + \mathbf{b}}$ = incidence for a 62 year old

7 SOA Spring Meeting, Anaheim
Probability of LTC Ruin? Session 21
Dan Cathcart, Amy Pahl, Janet Perrie
May 19, 2004

Chart 2

Severity

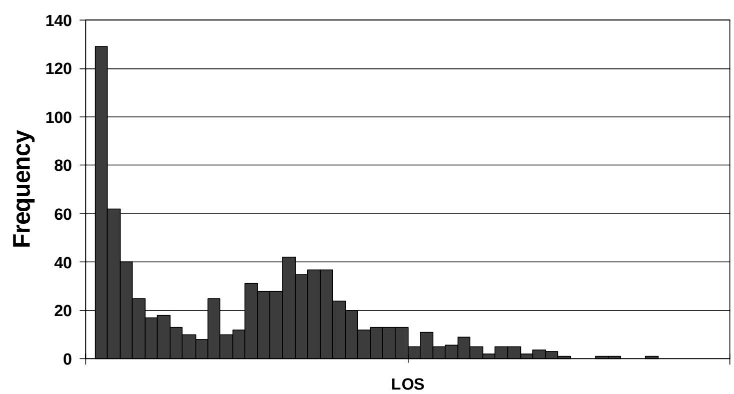

Mean = 1,171 days Std Dev = 955 days

10 SOA Spring Meeting, Anaheim
Probability of LTC Ruin? Session 21
Dan Cathcart, Amy Pahl, Janet Perrie
May 19, 2004

Chart 3

Deterministic Pricing Model

Annual Premium = $3,734

Component	PV* as a % CP
Collected Premium (CP)	100.0%
Premium Refunded	(1.4)%
Investment Income	42.6%
Claims Paid	(55.1)%
Reserve Increase	(39.2)%
Expenses	(24.8)%
Tax	(9.2)%
Change in Target Surplus	(3.2)%
After Tax Profit (15% IRR)	9.7%

* All values discounted at 5.5%

Chart 4

Stochastic Results – 10,000 Trials

Mean = 9.9%