RECORD, Volume 22, No. 2*

Colorado Springs Meeting
June 26–28, 1996

Session 114TS
Measuring and Analyzing Volatility Risk in Disability Income

Track: Education and Research/Financial Reporting
Key words: Disability Insurance, Reserves, Valuation of Liabilities

Co-Instructors: JEYARAJ VADIVELOO
DARRYL G. WAGNER
Recorder: RICHARD M. RASIEJ

Summary: Inability to qualify the volatility in individual disability income (DI) insurance is one of the key reasons for loss experiences in the industry. Product design and pricing have failed to reflect this risk. This session provides:

- *a volatility measure to determine reserves adequate to cover the statistical fluctuation in morbidity,*
- *an application of the reserve adequacy criteria under various plan designs,*
- *an application of the volatility measure to develop risk-adjusted premiums, and*
- *a test of the effectiveness of the volatility measure by simulating a block of business and computing ruin probabilities and the impact on surplus.*

Mr. Darryl G. Wagner: I'll be both your moderator and one of the presenters. Jay Vadiveloo is the head of reserving for Aetna Retirement Services in Hartford, Connecticut. He recently came to the Aetna from Connecticut Mutual where he worked for 13 years. Among his responsibilities at CML as appointed actuary, he oversaw a significant block of individual disability income business. He's also active as an adjunct professor at the University of Connecticut, and it's really through that connection and working with a graduate student there that he developed the methodology that we'll be talking to you about.

I am a consulting actuary with Arthur Andersen also in Hartford Connecticut, and a considerable part of what I do involves working with clients on a variety of issues related to both individual and group disability products.

Our objective is to introduce you to a methodology for measuring the volatility risk associated with disability income. The value of this methodology, we believe, is that it facilitates both the anticipation and the quantification of volatility risk in disability business which, if used in that way, allows you to better manage this risk. The old adage of it's hard to manage what you can't measure very much applies to the volatility risk with disability income, and that's the focus of what we're going to talk about. I'm going to start things off and really set the stage in terms of why this sort of measure is important. Just what do we mean by volatility? Why is it that Disability Income (DI) is volatile? Why is this type of methodology important?

Jay will pick up from there and present the methodology itself, giving a description of it and some examples of how it works in practice. Finally, I'll wrap up and discuss some of the conclusions and the impact of the methodology.

In my definition of volatility risk, we focus on the statistical fluctuation of results around some expected level, which is to say we're not really talking about misstatement or pricing risk. The assumption is you know over the long-term and in aggregate exactly what your claims experience is going to be for disability, but what we're talking about is the day-to-day, month-to-month, year-to-year, and policy-to-policy fluctuation that we all know occurs around that and how we can better measure that. It does not include shifts in underlying morbidity. I think there have been many sessions in the last few years that have talked about shifts in morbidity and general increases in claims cost and what they're attributable to. We're going to touch on that briefly, but it's really not the focus of this presentation.

As I said, we're going to focus on volatility risk, and specifically we're only going to talk about volatility with respect to morbidity. There certainly are many other causes of fluctuation in results for disability—interest rates, asset/liability matches, and so on. It's not to say that these things are not important, but to really focus on this methodology and presentation, we're going to focus on the morbidity risk. We're also going to be showing you something specific to individual disability insurance, but I think many of the concepts could carry over to the group disability product, as well as other types of insurance products.

I'd like to talk about the current DI environment. The pricing methodology that's used for DI products, in a broad sense, is comparable to that which is used in life insurance products; that is to say, were projecting some sort of expected claims cost and with some kind of underlying modeling to come up with a pricing structure. We'll talk about some issues around that in a minute. Also, with the individual product you typically have a premium level that's guaranteed through the life of the contract, or the noncancellable product design is common. There's also a

multiplicity of contract variations: benefit periods, elimination periods, types of coverages, amounts of benefits. There are many other bells and whistles out there.

Another characterization of the current environment is that there are adverse results. I reviewed a few headlines in getting ready for this presentation, and one of the headlines was "DI Insurers Try to Stop the Bleeding" which certainly says a lot just in and of itself. There certainly have been many problems, particularly financial losses. I think you can attribute that to two reasons. One is an increase in claim costs overall, but I think this issue of volatility is also a big part of that.

Is the claim cost at the right level to start with? I think there's a number of things that have driven this increase in claim costs. There has been much discussion of this in the industry. I think the main reason is the failure of risk classification. A favorite example of this is physicians who, for various reasons, have exhibited much higher-than-expected claims cost and have caused problems for many DI insurers.

Another factor is the richness of the benefit offerings. I was in another session at which the speaker said many of the benefit offerings that were developed in the 1980s could really be described as experimental. We weren't really sure what they were going to lead to and how we were going to price for them. I think we've seen that the richness of the benefit offerings themselves, along with a weak economy, have fueled subjective claims decision. Whether or not you're on disability has a fair amount of subjectivity to it and I think that is a risk in this area. There are also new causes of disability, such as AIDS and HIV. Inadequate claims administration and intervention have been problems in the past, and I think efforts are being made in all these areas to improve, but these certainly have been contributing factors.

Let's move on to volatility, which again is where we want to focus. Why is DI volatile? What is it about DI that leads to volatility? A few thoughts on that. First, there are the small and varied pricing classes. We talked about that before. There's a multiplicity of designs, which in and of itself means you're taking a large group of insureds and funneling them down into several small groups. I think this is one of the reasons why the application of a "life insurance type pricing methodology" can lead to problems because the spread of risk may not be the same. Also, there's variance in not only the incidence of claims but also the severity; in other words, not only are we projecting how many claims, but for each claim that occurs we're projecting how long it is going to endure. This is more of a casualty-type risk in many ways than what we see in the life insurance industry. Also, there's the extended duration of the claims themselves. You can easily have 30–40-year claim durations. There's a relatively low number of claims related to the number of insureds, and that, again, is kind of a spread-of-risk issue. Finally, there's the subjectivity regarding claim status. If you're selling life insurance, it's not too

subjective to say whether someone is alive or dead, but with disability it's certainly not that objective.

What has been the reaction to these various issues? First, many companies have gotten out of the individual DI business. There also have been switches to a guaranteed renewable product design, giving the flexibility to change premiums in the future to a greater extent than under noncancellable kind of coverage. Also, benefit restrictions and limitations have been added. A particular example of that is disabilities attributable to mental or nervous conditions. Also, sex-distinct pricing and state pricing differentials have been used to better focus in on risks in those areas. One problem, though, with that is that it adds to the multiplicity of product variations and these small buckets of insureds.

There is tighter underwriting and claims management. Think of the example of physician groups. Certainly companies have made changes there, as well as to their overall underwriting. This is a strategy, I think, that has been used by some DI carriers, in effect, giving them that larger spread of risk through taking on some of these large groups. Finally, there's restructuring, which I think more often than not translates into increasing the prices of product offerings. This has been a common approach.

In all of these reactions, however, the fundamental problem of how you measure, analyze, and price for fluctuation volatility in morbidity experience is really not addressed. I think that sets the stage for what we're going to talk about with this methodology. Even if we tighten the underwriting and figure out where the claims costs are, we still have this issue of volatility. An analogy that comes to mind for this is, "I got the short end of the stick." Well, it's one thing to say that, but it's another thing to not know the stick had a short end or to not know how short the end of that stick is. I think that's what we really want to talk about here. That statistical fluctuation is always going to be there, and there are going to be times where you get the short end of the stick. How can we be better prepared for that?

Mr. Jeyaraj Vadiveloo: Let me give you a description of the outline I'll be following. I will be defining volatility risk, and I'll be defining it in terms of claim reserves. With this definition I shall develop what are called risk-adjusted claim reserves or adequate claim reserves, and from risk-adjusted claim reserves I shall develop risk-adjusted claim costs and risk-adjusted active life reserves. Finally I shall develop risk-adjusted gross premiums which is a required step-by-step process. I will conclude my portion of the presentation by describing an approximation method or a manual algorithm which will do all of the above but sort of circumvent the simulation process.

Before I get into the details of my presentation, I'd like to make a couple of comments. All the results that you'll be seeing are the work of a graduate student at the University of Connecticut. His name is Hong Dai. In fact, this work is part of his Ph.D. dissertation, and I'm helping out as one of his thesis advisors. Hong Dai is not here, but I do want to acknowledge and thank him for all his work. Before I left Connecticut, I received some really exciting news. I just found out that this research has been awarded a Society of Actuaries' grant, and that's a real honor for Hong Dai and the University of Connecticut. I do want to thank the Society for recognizing this research.

I'm not sure how many of you subscribe to the *Disability Income Newsletter* put out by Milliman & Robertson, but the most recent newsletter talked about statistical fluctuation in morbidity, and that's similar to the concept we'll be discussing. I'll be looking at very different measures and results. This approximation method I'll be talking about is also a new idea, but to me the good news is that finally attention is being drawn to this quantification of volatility. Everyone is aware of volatility in DI. The problem is how do you quantify it? As these ideas get entrenched, it will change the way the DI business is going to be run in the future.

Let me define adequacy of claim reserves. I shall use sort of a standard probability type approach. I'll define the reserves as adequate if there's a strong likelihood that all your claim payment obligations can be met. I will quantify *strong likelihood* by attaching a probability number to it—a one minus alpha probability criterion. If alpha equals 10%, that'll be a 90% probability criterion or adequacy criterion.

The model is a basic Monte Carlo simulation. You look at a closed block of DI claims, and you use experience, monthly termination rates, and each simulation projects the claim or the block of claims to the end of the benefit period. In the classic Monte Carlo approach, you repeat the simulations. We did several thousand simulations per set and, of course, each simulation is equally likely. The measure we were capturing was the present value of monthly claim payments. You plot a probability distribution, and the one minus alpha percentile point is precisely your risk-adjusted claim reserve. So, if alpha is 10%, you'll choose the 90th percentile point.

In all the results I'll be showing, I shall use ten assumptions on the model. There will be some special situations where I'll be using those different assumptions, but it'll be clear from the examples when I use different assumptions. In general I'll be using a 90-day waiting period, a 4A occupation class, or professional class. I shall stay with a $1 monthly benefit to age 65. For valuation assumptions, I shall use the 1985 Commissioners Individual Disability Table A (CIDA) at a 5% interest rate. For modeling purposes, I'll be setting experience assumptions to equal valuation

assumptions. Obviously, you don't have to do that, but it makes the modeling easier to explain. I'll also use a 90% adequacy criterion.

In all the results that I'll be showing you, you'll see two measures that I'll be capturing. One is the claim reserve adequacy factor. I also call it the reserve strengthening factor. And the other factor is the percentage loss under ruin. What are these measures? The claim reserve adequacy factor is simply your risk-adjusted reserve level divided by your valuation reserves, minus one, and expressed as a percentage. For instance, if your claim reserve adequacy factor is 20%, and your tabular reserves are $100 million, to achieve your risk-adjusted claim reserve level, you should increase your tabular reserves to $120 million.

The second measure is the percentage loss under ruin. I shared portions of this presentation with some regulators, and one of the questions they raised was, it's fine to come up with an adequate level of reserves, but what happens in the event that you fail (even though it's a small probability)? What is your ruin cost? We decided to capture a conditional expectation number. Given that you have been ruined, what's the expected loss? The percentage loss equals your average loss expressed as a percentage of your risk-adjusted claim reserve.

To carry on the example, suppose your tabular claim reserves were $100 million, and you had a 20% claim reserve adequacy factor. You would increase your reserves to $120 million. Suppose you did 1,000 simulations. If you had a 90% criterion, about 100 of your simulations will fall in the tail. For every tail observation you look at the present value of claim payments and subtract the present value of claim payments at the 90th percentile—that's your loss. Then you compute the average loss among these 100 observations. Suppose the average loss is $12 million. Your risk-adjusted claim reserve level is $120 million. Your percentage loss under ruin is 10%. It's $12 million over $120 million.

We are ready to get some results. In Table 1, I was trying to determine the effect of different incurral ages of disability. So, we looked at incurral ages of disability at 25, 35, 45, and 55, and I crossed it against claim duration. So zero is a new claim, 12 months is a one-year-old claim, and 24 months is a two-year-old claim. If you look at incurral age 45, for a new claim, you see a claim reserve adequacy factor of 12%. That means that your tabular reserves have to be multiplied by 1.12 to come up with your 90% adequacy level. However, for that same incurral age, by the time you reach a two-year-old claim, your claim reserve adequacy factor is only 4%. You'll see this pattern that shows for the younger incurral ages and the early duration claims, you have the greatest volatility. As you move to the older incurral ages and the longer duration claims, your volatility factor and reserve strengthening

factor start to decrease. When you look at the percentage loss under ruin (Table 2), you get the same pattern, but the factors are all basically smaller in magnitude.

TABLE 1
IMPACT OF INCURRAL AGE OF DISABILITY
CLAIM RESERVE ADEQUACY FACTOR

Incurral age of disability	Duration from end of waiting period (in months)		
	0	12	24
25	15%	14%	9%
35	14	11	6
45	12	8	4
55	8	6	3

TABLE 2
IMPACT OF INCURRAL AGE OF DISABILITY
PERCENTAGE OF LOSS UNDER RUIN

Incurral age of dis-ability	Duration from end of waiting period (in months)		
	0	12	24
25	8%	4%	3%
35	7	3	2
45	5	2	2
55	3	1	1

One nice thing about looking at both these factors is it has always been an issue, when you are coming up with an adequacy criterion, to determine the alpha level. Should it be 10% or 5% or 1%? It's always somewhat of an arbitrary decision. However, if you combine this percentage-under-ruin factor with the claim reserve adequacy factor, the whole debate on what your level of confidence is becomes a moot issue. There is one way of using this combination of results. Suppose again, for an incurral age 45, for a new claim, you sort of do a two-step procedure. You first increase your tabular reserves by 12% because that's your claim reserve adequacy factor for a new claim. Then, to cover the ruin probability you further increase it by 5% because that is the percentage loss under ruin factor for a new claim at age 45.

What happens when you look at the changes in your waiting period? In Table 3, you get a pattern that you might be anticipating. The shorter waiting periods have high volatility. Again, you still get the pattern at the older claim durations that shows the volatility going down. For instance, a new claim with a 30-day wait shows that the volatility factor is 17%, and for a 180-day wait, it's only 8%. When

you go to the percentage loss under ruin (Table 4), you get the same pattern, again with the numbers significantly smaller.

TABLE 3
IMPACT OF WAITING PERIOD
CLAIM RESERVE ADEQUACY FACTOR

Duration from end of waiting period (in months)	Waiting period (in days)			
	30	60	90	180
0	17%	15%	12%	8%
12	10	9	8	6
24	4	4	4	4

TABLE 4
IMPACT OF WAITING PERIOD
PERCENTAGE LOSS UNDER RUIN

Duration from end of waiting period (in months)	Waiting period (in days)			
	30	60	90	180
0	8%	7%	5%	4%
12	2	2	2	2
24	1	2	2	2

When we moved onto the impact of the length of benefit period, the results surprised us because we saw this big jump in volatility risk moving from a limited benefit period, such as one year, two years, and five years, to a benefit period to age 65 (Table 5). The longer benefit periods, as expected, have high volatility. At the older duration claims, you have fairly negligible volatility and basically a similar pattern for the percentage loss under ruin (Table 6).

TABLE 5
IMPACT OF LENGTH OF BENEFIT PERIOD
CLAIM RESERVE ADEQUACY FACTOR

Duration from end of waiting period (in months)	Benefit Period (in years)			
	1	2	5	To age 65
0	1%	1%	1%	12%
12	—	6	6	9
24	—	—	3	4

TABLE 6
IMPACT OF LENGTH OF BENEFIT PERIOD
PERCENTAGE OF LOSS UNDER RUIN

Duration from end of waiting period (in months)	Benefit Period (in years)			
	1	2	5	to age 65
0	2%	3%	4%	5%
12	—	1	1	2
24	—	—	1	2

One of the general observations that we can make is that when benefit provisions are liberalized, and by that I mean you have shorter waiting periods and longer benefit periods, you have greater volatility. As I said, that is somewhat intuitive. What may not be so intuitive is that at the older incurral ages and at the longer duration claims you have little volatility. Remember, we are not measuring the claim reserve yet, but we are measuring the volatility of the claim reserve. The claim reserves are higher at the older durations and at the older incurral ages, but because termination rates are so small, there's very little fluctuation.

With regards to the percentage loss under ruin, as I mentioned, it has the same pattern as the claim reserve adequacy factor except that the numbers are significantly smaller. It gets pretty negligible at the older incurral ages and the longer duration claims.

We did many tests on this next observation. We tried to shift the morbidity table. Instead of using 1985 CIDA, we multiplied all termination rates by 90%, but still capped experience assumptions equal to valuation assumptions. The claim reserve adequacy factors and the percentage loss under ruin didn't change significantly. So it's quite invariant under shifts in your underlying morbidity, and it also had the same lack of change when you changed the discount rate. Again, you had to set experience assumptions equal to valuation assumptions.

The next observation wasn't so obvious. What happens when you have a mixed block of claims and you vary the volume as well as the distribution of claims? What does that do to the overall claim reserve adequacy factor? Well, in general, as the number of policies on claim increases, your volatility decreases. That's basically the law of large numbers at work, so that's to be expected. However, the second observation was, it was a very nice result. I didn't expect it to be so nice. When you have a mix of policies your reserve strengthening factor turns out to be a weighted linear combination of the factors for the individual groups. There are a couple of important caveats, and I want to mention them. These factors for the individual groups are not based on the count within these groups. They are based

on the total claim count, and that's very important. The weights are not based on count but are based on the monthly benefit amount.

Let's discuss a simple example. Suppose we have 360 claims, all with a 90-day wait, and all fall under the 4A male. Of those 360, 180 are new claims and 180 are two-year-old claims. We have the same monthly benefit for all claims. The claim reserve adequacy factor for new claims is 12%, and for two-year-old claims it's 4%. Both this 12% and 4% factor is based on the total count of 360 and not the individual counts of 180. So the combined or the weighted reserve strengthening factor is just a simple average because you have equal monthly benefit amounts for both blocks. It's half of 12% plus half of 4%, and that gives you an 8% weighted average. If I doubled the monthly benefit amount for the new claims, then my weights will change. It'll be a two-third weight for the 12%, and a one-third weight for the 4% group, and that gives you an overall claim reserve adequacy factor of 9.33%.

All the simulation work stops at the claim reserve level. When you build up the risk-adjusted claim costs and active life reserves, you just go backward in a systematic fashion. This is the step-by-step process. First you use the simulation method to come up with your risk-adjusted claim reserves, and then the next step is to come up with your risk-adjusted claim cost, and your risk-adjusted claim cost is simply your experience incidence rate multiplied by your risk-adjusted claim reserve for a new claim. Then, once you get your risk-adjusted claim cost, you use this risk-adjusted claim cost and develop your risk-adjusted active life reserves the usual way.

When you go through this sort of systematic approach, you'll find the factors for the active life reserves will be significantly smaller than the factors for the claim reserves. The other observation is that when I came up with the risk-adjusted claim cost, I used the experience incidence rates and then multiplied it by the risk-adjusted claim reserve. I didn't use a risk-adjusted experience incidence rate. The reason you don't have to is the law of large numbers. All healthy lives are exposed to the incidence rate, and the incidence rate is a binomial random variable. You either get on claim or you don't get on claim. And the convergence of your experience incidence rate to your true underlying incidence rate is fairly rapid. You don't have to use a volatility envelope around the incidence rate.

Once you get your active life reserves, your risk-adjusted active life reserves, and your risk-adjusted claim cost, you develop your risk-adjusted gross premiums in a two-stage process. You first develop what I call risk-adjusted benefit premiums, and they just satisfy the usual actuarial formula. The actuarial present value of risk-adjusted benefit premiums equals the actuarial present value of risk-adjusted claim

cost. And from your risk-adjusted benefit premiums you come up to your risk-adjusted gross premiums by just adding the loading for profits and expenses. That's the basic process. You'll find that by the time you arrive at the risk-adjusted gross premiums the factors are, again, significantly smaller than the starting factors for the claim reserves.

We used, as I said, a full-blown simulation to come up with the 90th percentile for the claim reserves, and then went through this backward procedure to come up with risk-adjusted gross premiums. It wasn't clear to us that your gross premiums had a 90% criterion within themselves because we didn't use the simulations for gross premiums. We tested it, and in general we found that the gross premiums were conservative. So, if you satisfied a 90% criterion for your claim reserves, by the time you went backwards to your gross premiums it was possibly a 95% adequacy level. It looked good, and it was conservative.

So we did all the simulations, and I was fortunate to have had access to all of the University of Connecticut's computing facilities. For a company to actually manage its disability income business using this method you would want to incorporate this process into your reserving, and, of course, if you do your valuation once a month, like most companies do, you cannot be doing the simulation process. It takes a great deal of work and effort. You could do it, but it can be quite cumbersome.

Hong Dai and I did much creative thinking to come up with an alternative method where we could bypass this whole simulation process. That was possibly the most interesting part of this research. None of us knew what this creative idea would lead to. What was nice about this manual algorithm, when we finally came up with it, was obviously you didn't have to do the simulations, so it was a good method. It was actually quite easy to do, and the really nice thing about it is it is actuarially very intuitive. That's what I liked about it. I believe it's so easy that you can build it right into your reserve systems and routinely calculate your risk-adjusted claim reserves and risk-adjusted active life reserves as you do your valuation. Finally, it is a good approximation. We did a lot of stress tests where we compared this approximation method to a full-blown simulation method, and, as I said, it is a pretty good approximation.

Let me describe this manual algorithm. What do you do? It's a seriatim process. You do it for each claim. Start with a policy on claim and do the following. You first have to calculate what I call a natural claim reserve, and the natural claim reserve is a regular reserve where you're using experience termination rates and a realistic earnings rate for discounting. Next, you come up with a one minus alpha adequacy level so you can use an alpha equal to 10% and come up with a 90% adequacy level.

The first step is to determine what I call the *r* factor. You look at the natural claim reserve that you have calculated and determine how many monthly claim payments can be sustained by this natural claim reserve. Suppose *r* is 30 months. That means if this policy stayed on claim beyond 30 months, your natural claim reserve would be inadequate. If the policy terminated within 30 months, your natural claim reserve would be adequate. So, that's the first term you determine. Then using your one minus alpha criterion, say 90%, and using your experience termination rates, you determine the second term—the *s* factor. You calculate how many monthly payments you need to cover so that you are at least 90% certain of being adequate? For instance, if your *s* factor turns out to be 80 months, that means there is at least a 90% chance that this person will terminate within 80 months. Another way of looking at it is there's less than a 10% chance that this policy on claim will persist beyond 80 months. Again, that's something you can calculate.

Two points. If you want to build this into your reserving systems, you must have some rules for the extreme cases. One extreme case would be, what happens if the probability of making it to the end of the benefit period exceeds 10%? You would just set your *s* term to equal the number of months to the end of the benefit period. That's just a rule. You have to set some rules for the extreme cases. The second point is an important point. When you're calculating *r* and *s,* you're using interest rather than interest and survivorship. It's not a difficult calculation.

So, what then is your risk-adjusted claim reserve using this manual method? Your risk-adjusted claim reserve is the sum of two reserve pieces. It's your natural claim reserve plus a constant (and it's an important constant), times a "deferred" reserve. It's a deferred, temporary, disabled, life annuity reserve after *r* claim payments and until *s* claim payments. I call that second piece your risk claim reserve. Effectively your risk-adjusted claim reserve is your natural claim reserve plus this constant multiplied by your risk claim reserve. As I said, it's very actuarial. What you're doing is you're saying that your natural claim reserve covers you up to a certain number of monthly payments, but you want to reach your adequacy level which requires you to go further. For that extra reserve, you just hold a deferred reserve. It's not additive, though. You have to multiply by this constant factor. There's a big difference.

If *r* is 30 months (your claim reserve sustains you for 30 months only), to hit a 90% adequacy level you have to go up to 80 months. Then your risk claim reserve portion is the actuarial present value of the 31st monthly claim payment until the 80th monthly claim payment.

Again, if you're going to build this into your valuation system, you have to make some rules for the extreme cases. In the rare instance that *r* exceeds *s*, your natural

claim reserve covers your adequacy criterion. The simple rule is just set your risk-adjusted reserve to equal your natural claim reserve. A little hint about the factor *c* which is still part of the research of my student is that it's actually a very well-behaved constant. It varies somewhat linearly (with the number of policies on claim). The more policies you have in your books, the smaller the value of *c*. I said it's a very well-behaved constant. I hope what I've described is not too complicated. When you actually have to work with it, you can very easily program this procedure right into your reserve systems, and you can do it on a seriatim basis. Now, if you didn't want to do that, and you want to do it on a group basis, you can really come up with your risk-adjusted claim reserves using an Excel spreadsheet. Once you get your risk-adjusted claim reserves, you do the rest of the processes the usual way because they don't involve simulation in any case. You do your risk-adjusted claim cost, your risk-adjusted active life reserves, and finally your risk-adjusted gross premiums. That is a step-by-step procedure on how to capture volatility, quantify it, and build it right into your reserving and pricing.

Before I turn the presentation back over to Darryl, I want to leave you with one thought for the day. It's a quote from Mark Twain, and I think it's very apt for this presentation. Mark Twain used to say that there are two instances in life when one should not speculate—first, when one cannot afford to and, second, when one can. That's how I view writing DI business without knowing how to measure and manage volatility. Whether you're a successful company or a struggling company, if you don't manage volatility, DI business just becomes a speculative venture. I hope that what I've presented here will give you some practical tools on how to measure and manage volatility.

Mr. Wagner: I'd like to talk about the conclusions that we've drawn from this and how the they might be applied within the context of managing a disability block of business.

Just to recap, the methodology that Jay presented gives you a measure of volatility risk, and I think that's the key output of this methodology. Another way to think of the claim reserve adequacy factor that Jay described, is really as a measure of volatility. We looked at those charts showing where there's a higher claim reserve adequacy factor and where there's a lower one; I think another way to look at that is to say where do I have higher volatility and where do I have lower volatility? What we're saying is that's what's needed to get you to a certain confidence level on a statistical basis.

I'd like to talk, from a couple of perspectives, about how this might be used. First is from the context of a company managing a block of business, and another is from a regulatory point of view. First, what does this do for a company? I think there are a

number of things. Basically, it allows a company to more effectively manage the DI business, specifically the volatility risk. How does that happen? First, it gives you a way to make the different levels of volatility that are inherent in different product designs a little more concrete so you can measure them. With that information in hand, it allows you to perhaps focus your product design and your pricing a little better. Then you can determine what your risk tolerance is. We're making a decision between a couple of benefit designs or a couple of marketplaces. This is another measure or another piece of information that you might add to your process. It is in addition to whatever information you have in place today. If I'm making a decision between two things, knowing what the volatility is may not affect my decision, but it may affect how I carry it out.

Let's discuss other applications. Once the decision is made to go into a particular product or market, if you're using something like reinsurance, you have a measurement to come back to if you structure a reinsurance agreement in a particular way. You can determine how it affects this volatility measure and, in some sense, it gives you a measure of success in terms of transferring risk or hedging risk. Traditionally, risky markets could be penetrated if this volatility could be reduced. I think what we're asking is, what constitutes a risky market? This volatility risk may have a lot to do with quantifying that. Maybe I modify my product design or maybe I bring in reinsurance. I have a measure for how effective those things are. The final example is a risk premium experience refund mechanism. Perhaps I'm using this volatility measure to increase prices on more volatile blocks of business. There could be some sort of release mechanism in conjunction with that, such that, over time, if that proves to be an unneeded margin, that could be released to policyholders in some way. We certainly haven't worked out all the details of that kind of a mechanism, but it's doable in principle.

I don't know if we have any regulators in attendance. From a regulatory perspective disability income is a challenge just as it is from a company perspective because of the same types of uncertainty and volatility that we've talked about, and I think one approach for this would be perhaps to link minimum reserve standards to this methodology. As I've worked with different companies, I've heard a great deal of discussion that the disability reserve assumptions, morbidity assumptions in particular, really don't do the job on both an individual and a group basis. What we're suggesting is coming back to experience-type assumptions but linking to them this methodology that would give us some sort of measure of the volatility. It goes back to that adequacy level that Jay talked about. I think this would be more of a move toward the valuation actuary type concept, and it would be a fairly radical change. Again, it's one I think that's worth considering.

This methodology might be something that could be used to demonstrate adequacy, either within a cash-flow testing or some other context. I think it could have applications into risk-based capital requirements, as well as dynamic financial solvency. What we're talking about is the volatility on the liability side of this business. When we do cash-flow testing and dynamic solvency analysis, there's often more of a focus on the volatility on the asset side of the books and perhaps not as much on the liability side. What we've presented here really would round that out.

As I sat in another session, I thought of perhaps another application of this methodology, which would be to the determination of the fair value of liabilities. As most of you probably know, under GAAP accounting, you now have requirements to mark your assets to a fair value or market value, in many cases. What needs to be developed and what has lagged is coming up with a methodology to do that same thing for the liability side of the balance sheet. I think this methodology presents some possibilities for that, specifically as it applies to disability income.

Finally, in this presentation, we didn't talk about other business risks, such as expense, asset risks, and even the overall misstatement of morbidity risk. We are not saying that these are not important. Certainly any consideration of reserve adequacy or pricing adequacy should consider these risks, in addition to volatility risk. As we've said, the focus here is on volatility. Finally, we're not presenting this as a panacea that takes the place of other practices that you might have in place such as stricter underwriting, good claims management, and benefit restrictions. We would certainly advocate continuing all these things. What we're saying is that the combination of those types of practices and addressing the volatility risk is going to be vital to success in this business as we go forward. That concludes our prepared remarks, and we'd certainly be happy to answer any questions that any of you may have.

From the Floor: I have a question for Jay. It wasn't clear to me that your manual algorithm was applicable to claims with a lifetime benefit period. Can it be modified to work with lifetime claims?

Mr. Vadiveloo: Yes. You can't easily do it because it isn't strictly additive which was a real surprise to us. That constant factor is critical. If you add the two pieces, it will just be too much, and that constant which varies the total number of claims, is a big factor. I am sure we can do modeling work to figure out a similar relationship for lifetime benefits. My student will be very happy to continue his research.

Mr. Richard Noel Ferree: Regarding those adequacy factors that you presented earlier—were they for a single claim or 100 claims?

Mr. Vadiveloo: It was for 360 claims. When I did that example, the 12% and 4% came precisely from those tables as well. What's interesting is we ran it all the way up to 1,500 claims, and you still have a nonzero claim adequacy factor. So the convergence to zero volatility takes a while, and that was also an interesting observation.

From the Floor: For 1,500 claims, was that about four times 360, so about half the volatility then?

Mr. Vadiveloo: I don't have the numbers in mind. It's not as simple because it varies depending on the section's durational claims and things like that. It does come out quite a bit, but then it stays fairly unchanged for a while. So it almost reaches a minimum and stays there for a while. We did the test up to much more than, I think, 2,000 or so claims. We were taking a great deal of computer time by then.

Mr. Lawrence N. Segal: Jay, it sounds like you're doing a lot of modeling to determine what *c* ought to be. Is it the case that if you are using this manual algorithm you'll still need to do quite a bit of modeling? Wouldn't the company still need to do a lot of modeling to figure out what is the appropriate *c*?

Mr. Vadiveloo: What I'm hoping is that, in fact, when the student publishes his work, he will specifically define the *c* factor. It's very stable, as I mentioned. I don't think you will have to, as a company, keep doing simulations to figure out *c*. You can quantify many of the situations this way.

Mr. Edwin H. Betz: I was a little bit confused when you indicated that if you combine two blocks of business, that the reserve adequacy factor would be a linear combination of those. Maybe my intuition is wrong, but I would have expected it to be lower than that because of the larger size of the pool.

Mr. Vadiveloo: Yes, that's why I said it's a weighted linear combination, but it uses those factors for the individual groups based on the total number of claims rather than the claims in the individual sections. Think back to that example that I did with a total of 360 claims, and then 180 in each. When it came up with the individual factors I used the count of 360 not 180. So I think that is including the overall size.

Mr. Wagner: Just to add to that, if you have a larger block of business, the individual factors themselves will be lower before you apply your weighted average.

From the Floor: When adding two blocks of business, you have to calculate the reserve adequacy factors for the two pools as though they were already a single pool? Is that what you're saying?

Mr. Vadiveloo: Yes, basically.

Mr. David Morgan Andreae: Have you given any consideration to residual benefits? Would the adequacy factor for residual payments vary from 25% to 100% of the factor for full benefits?

Mr. Vadiveloo: There are many issues like this. Ultimately, I think when you want to come up with a manual algorithm, you will have to make certain rules. I think these are all very good questions for me to pass onto my student when I go back—these are things we'll want to test—how sensitive is it for partial benefits or residual benefits. We'll certainly be testing things like that.

Mr. Robert G. Meilander: The question I have for you is, why 90%? How does it relate to surplus requirements under risk-based capital analysis?

Mr. Vadiveloo: No special reason. I told you that if you also capture the percentage loss under ruin and build that into your ultimate reserve strengthening, it doesn't really matter so much whether you choose 90% or a 95% or 99%. We just picked a number and stayed with it.

From the Floor: It seems that a portion of that is actually a risk-based capital number, and, unfortunately, that's the reserve number. I would agree with the total. That's probably the right total.

Mr. Vadiveloo: Yes.

From the Floor: You start with a different adequacy criterion?

Mr. Vadiveloo: Certainly the manual algorithm can easily be adjusted if you change your adequacy criterion from 90% to 95%.

Mr. Ferree: Did you test the ruin level of just the regular natural reserve to see if it is the 50th or 55th or 45th percentile?

Mr. Vadiveloo: Something tells you you can go backwards and solve for it from the numbers I've given, but we didn't. We did the simulations and looked at all possible present values of claim payments, and your natural claim reserve, I believe, wasn't at 50%, as people think. It's actually closer to a percentage in the high 70s.

It takes a great many claim payments to go from the high 70s or 80s to 90%. The tails are really long. At least I had all that information. All that information is available. I can always calculate the ratio of the ruin in terms of your natural claim reserve. We have that all in our database. I still think you can go backwards and solve for it from the numbers I've given.

Mr. Joseph Frank Talarico: Jay, have you done any testing to determine how sensitive the analysis would be if you did not use the experience assumptions equal to the valuation assumptions?

Mr. Vadiveloo: Well, suppose your experience assumptions aren't equal to your valuation assumption; typically they are not. In a way, it doesn't really matter. You always start out with your experience assumptions because that's the only real set of numbers you're working with. Then you calculate this risk-adjusted claim reserve. It's going to be some number. You represent it as a ratio of your tabular reserves. What you capture with this manual algorithm or even the simulation procedure, if you use your experience assumptions, is your risk-adjusted reserve.

Mr. Wagner: Just an added thought to that. You can think of this in a relative sense and in an absolute sense. In a relative sense, if you're just looking at a volatility measure, in some sense, it doesn't really matter what base you're starting from or what adequacy criterion you're using. You're looking at relative measures of volatility. If you're using it in an absolute sense to either add to premiums or add to reserves, then I think you put those under a much closer spotlight. However, as Jay said, I think you can basically start from any point as long as you're consistent in carrying that through.

Mr. Vadiveloo: Again, I do want to emphasize that you always have to start with experience assumptions because tabular assumptions are arbitrary, regulatory-driven numbers, and this whole volatility measure is volatility about your own experience.

Mr. Carl L. Loeffel: I assume that this does not change if you change the interest rate.

Mr. Vadiveloo: No, it doesn't change if I set experience assumptions equal to valuation assumptions. Suppose valuation interest rates were at 4%, and your earnings rate is at 8%. The risk factor is going to be quite different when measured against your tabular reserves, but the risk factor measured against valuation assumptions with an 8% interest rate will be pretty much the same.

From the Floor: I didn't mean to go from your valuation assumption, but suppose your valuation assumption was 0%.

Mr. Vadiveloo: Right.

From the Floor: Would you still come up with the same answer?

Mr. Vadiveloo: The same answer as in the factor or the same answer as in my risk-adjusted reserve? My risk-adjusted reserve will be independent of my valuation assumptions because it's based on experience assumptions. Once I calculate my risk-adjusted reserve, my 90th percentile, I'm going to take it as a ratio to something. You would typically want to take it as a ratio of your tabular reserves. So the risk-adjusted reserve will always be the same. The factor will change depending on what I take as a ratio. I showed in my tables that it's quite invariant to shifts in interest rates and morbidity if you set experience equal to valuation. It has some very nice, stable properties. That's what I liked about this analysis.

Mr. Ferree: It also looks like this could be applied to, say, long-term-care terminations or any kind of a life annuity.

Mr. Wagner: I think that's absolutely right, and I think the research was done specifically to disability, but I think that certainly for any extended claim-oriented policy like a long-term care or even group LTD, it could apply. Going beyond that, for instance, an annuity policy in some sense is not that different from a disability policy once someone has gone on a claim. I think there's applicability there, too. You may not feel there's much of a need for it there because of the amount of volatility risk, but I think there's an application.

Mr. Vadiveloo: Believe it or not, we are looking at another application, and it's on the life side. It's for second-to-die products, which have all the classic characteristics of high volatility because there are very low probabilities and very high benefit amounts. You have concentration risk, too. What I'm having to spend some time on is coming up with a manual algorithm. That's the hard part. I'm trying to be creative.

Mr. Segal: You said *r* and *s* were determined just with interest, but wouldn't the risk claim reserve reflect mortality as well?

Mr. Vadiveloo: Yes, absolutely.

From the Floor: Interest and mortality.

Mr. Vadiveloo: Yes.

From the Floor: I was surprised by your conclusions on the risk-adjustment factors, especially when you get out to 90-day and 180-day waiting periods. Second, I was also surprised by the interaction between a shift in an incidence assumption, an experience incidence rate, and the effect on the claim reserve volatility. I'd be interested in studying that and doing some analysis.

Mr. Vadiveloo: Yes. I totally agree that that's critical. These results only make sense if your underlying experience assumption rates are reasonable. If you have a misstatement risk, you don't measure your incidence rates correctly. Obviously, nothing I say holds. But remember, the only true sort of scientific approach was at the claim reserves. That's where we did a full simulation. Everything else was sort of a logical progression back. When I made that decision to use the experience incidence rates and then multiply by the risk-adjusted claim reserve, my conjecture was that you have so many healthy lives exposed to this incidence rate. It's a binomial random variable, and the binomial convergence is fairly rapid. You are actually capturing the volatility that you associate with incidence rates because you're multiplying by a risk-adjusted claim reserve. We independently tested the adequacy of the gross premiums, and we found that it was conservative. At least it worked when you finally did the test at the end.