RECORD, Volume 25, No. 1*

Atlanta Spring Meeting
May 24–25, 1999

Session 24PD
Balancing Risks

Track: Investment
Key Words: Asset/Liability Management, Financial Management, Risk Management

Moderator: NANCY E. BENNETT
Panelists: NANCY E. BENNETT
FRANCIS P. SABATINI
Recorder: PETER D. TILLEY

Summary: What is the approximate mix of financial risks? Should a company take more credit risk versus interest rate risk? How should a company manage the uncertainty of policyholder behavior (i.e., lapse risk)? Panelists discuss a comprehensive risk management approach to quantify the risk exposure from various risk contributors.

Ms. Nancy E. Bennett: The session is going to focus on the management of financial risks in an insurance organization. One central theme you're going to hear will be the importance of enterprise-wide risk management. Frank Sabatini and I are the presenters.

Frank Sabatini is very well known and respected in the actuarial community. He's a frequent speaker at SOA meetings. He has earned a well-deserved reputation as a leading expert in asset/liability and risk management. Frank is a partner in the Hartford office of Ernst and Young. I will be the second speaker. I have recently joined the Avon Consulting Group, where we are opening a new Midwest office. Prior to joining Avon Consulting, I worked for Minnesota Mutual Life Insurance Company, which has been renamed Minnesota Life. While at Minnesota Mutual, I was responsible for the company's corporate actuarial area, and was the company's appointed actuary. I have a lot of hands-on experience with cash-flow testing, integrated financial management, and risk management.

Mr. Francis P. Sabatini: We're going to talk about balancing risks, and we'll define that in a few minutes. It's bigger than a breadbasket, and it's broader than what we are used to referring to in terms of a risk management context. But first, let's talk a little bit about where we are. For most companies, the risk universe is interest-rate risk. We've been doing it for years. Some are better than others, and we understand it fairly well. With the growth in the variable annuity (VA) markets, the emergence of death-benefit guarantees, living-benefit guarantees on VA

Note: The chart referred to in the text can be found at the end of the manuscript.

products, and the emergence of equity-indexed annuities, there is increasing recognition of equity exposures, although they haven't yet filtered into the risk-management process. It is beginning to happen in many companies. You'll see the word "holistic" a number of times. And in today's environment, most risk management in most companies is not holistic.

And what does that mean? I'll tell you a story. I had a client who was concerned about the risk exposure that existed in a particular product line. There were a variety of factors on the asset side that pertained to the way the business was being run. They had done some risk analytics. The risk analytics bore out the risk concerns that the management team had, but it had all these other lines of business that didn't do any risk management. The team didn't know how to fit in the total. They couldn't make a decision because they didn't have a holistic view. They didn't realize that taking the risk exposures in the other product lines might have made the risk position of this particular product line relatively insignificant or less of a concern. In the few instances that I've experienced situations like that, the concerns frequently went away. If it's not holistic, it's not combined by product line or line of business. The key idea here is, holistic means that you can view risk across the entire enterprise.

Generally, the information that is being produced is produced more for internal use, by risk managers who frequently do some really neat things in terms of managing and measuring risk, but it has limited senior management emphasis. The information that's provided doesn't have tremendous utility. Most of the risk management that we do today has a long-term perspective. Ending surplus, present value of ending surplus, present value of distributable earnings, and other similar measures like duration and convexity are long-term-oriented measures.

Just in contrast, I thought it would be interesting to look at what has happened on the bank side. On the bank side, regulation has really forced a quicker evolution than there has been in the insurance industry. There's a broader view of risk, and you're likely to see holistic activity. Value at risk (VaR) is well established. It has its own faults, so I'm certainly not a 100% fan of VaR. It fits the bank model better than it would fit the insurance model. Banks look at all the risks. In some more sophisticated banks, you'll even see them factoring in the liabilities. Management is very much involved. VaR results go all the way up to the board level. What's interesting is the horizon is less than one year. That's primarily because their corporate securities are publicly traded securities on the balance sheet and marked to market. Price volatility of those assets is a very important part of the risk management program.

But I believe we're headed for transition in risk management, and I'd like to at least describe my vision of a future state. It's going to be the entire risk universe, and we'll spend a little more time on that. That risk universe could be anything from market conduct to litigation, to operational risk, to interest rate risk, to equity market exposures, or anything and everything. I think there are some companies that are beginning to implement the processes that focus on the entire risk universe. It'll be holistic. When you look at risk, you look at it across the entire organization, and you just do not stop at legal entities. For those of you who are

familiar with some of the more well-established risk management programs, one of the hallmarks is the fact that it's a very disciplined, very focused approach to managing and measuring risk. Of course, one of the hallmarks of most good risk management programs is that they actually produce information, and then the companies act on the information. It'll have tremendous management attention and utility, and we'll think of risk in more discreet time frames.

I have an example that will address that a little more. What's the impetus? Either the insurance industry is going to take over the banking industry or the banking industry is going to take over the insurance industry, or we're all going to be owned by foreign companies. We're in the process of demutualizing, which means that for some very large, new public companies, some of the issues that are surrounding risk management aren't as important as they will be once they're public entities. Competition is not getting any easier, despite all the consolidation. You would think that consolidation would mean less competitors, but it's getting more competitive. Shareholder demands and expectations never go away. That all leads to the need for broader and improved risk management tools and information.

I want to talk about comprehensive risk management. Here's a very simple, succinct definition. It's a risk and value measurement framework that incorporates all risks. Some of them may be really hard to measure, but at least they're part of the entire framework. There are some process-related items. It would be interesting if we could talk about all the different risks and develop a risk inventory. You'd be surprised at how long that list will be. Some of those risks will be fairly significant. I will argue that the risks we spend most of our time measuring and managing aren't the ones that are the greatest risks to our organization. First, we must identify what risks the organization has. That frequently involves sitting back and looking at all aspects of the business, and identifying all the different risk elements. A harder job is quantifying risk. Then you must try to find a way to aggregate it so you can achieve your holistic results.

Of course, once you have the whole perspective, then the challenge is deciding how much risk you're willing to absorb, or are willing to take, and how much appetite you have for risk. My experience is that there'll be as many risk tolerance levels in this room as there are people. In my experience with boards and senior management, when we worked on defining risk, we showed them the measures, they understood them, and then one said, "Well, that's okay." Another one will say, "Well, I think that's too much risk." And before you know it, you have 12 different opinions. Eventually there's some convergence, because they say what would it do to us if we went off the risk or took less risk? They say, "That's an awfully expensive price to take us off that risk." It might not be too bad after all, but you go through a very interesting process. Then you get into the real interesting stuff in terms of linking risks with returns. You start using the foundation you've built, the strategy formulation, and the capital allocation. You become a risk management process, and you actually manage risk, monitor it, and act on the information that you produce. Then you can get into things like performance measurement and attribution, and actually use the process to add value. Rather than managing risk, you're now managing to create value. In the end, the successful organizations have been able to institutionalize it.

I've taken the risk universe and tried to put everything into four basic categories. The first is financial, which includes interest rate risk, equity risk, lapse risk, policyholder behaviors, premium risk, production risk, and anything that has a financial nature to it. There is also business risk, which includes market conduct. Third is operational risks, which are, surprisingly important. Finally is event risk. All of these are the things that need to be considered, and part of an overall comprehensive risk management process. I'm not going to spend a lot of time on defining these, but they are important.

What are the benefits? One benefit is bringing all the risks together into a common measurement framework. This creates a perspective on relative exposure to various risks. Risk correlation is kind of a misnomer. Frequently, many of the risks aren't correlated, which means you benefit from the fact that they're not, or the fact that they're negatively correlated. That's something that we don't spend a lot of time on. The case study that we're going to cover is going to illustrate that whole issue in terms of the benefits of combining risks, the benefits of capitalizing on the fact that the risks are either independent or negatively rather than positively correlated. Finally, it provides a foundation for strategic decision making and other value-added activities.

I'm going to talk about methodology in the context of financial risk management, and in the context of some of the things that we're more familiar with. In my risk-management world, there's no such thing as straight application of constant mortality rates. There's no such thing as annualized default rates. If we're measuring and managing risk, why would we build a process that ignores the fact that there are mortality variations and trends in mortality? We can't ignore that. In my world, each risk element is typically defined by some statistical technique. It might have a mean, standard deviation, or it might have some other statistical basis. We're going to deal with specific and systematic risk. *Specific* means pure random fluctuations, from period to period. All of you, I'm sure, have seen that mortality is up this year. Why? I would argue it's pure random fluctuation. As for trends, the one thing that you can probably relate best to is mortality improvement. We're going to define the risk profile reflecting all elements collectively, and we'll recognize their relative correlation or lack of correlation. The beauty of the entire process is now you can start talking about the relative contribution of each risk element to your overall risk position.

Let's do a little survey. How many people here think they have more interest rate risk than credit risk? Interest rate is greater. How many people think there's more credit risk in their companies? Have you quantified it? It's hard to do, and that's where we're headed.

We're actually going to do financial risk management, which is a subset of comprehensive risk management. We're only dealing with financial risks. In some respects, it is a fairly simple case study, and in other respects, it is a fairly sophisticated one. We're going to deal with a typical universal life liability flexible premium with all the bells and whistles, like cost of insurance (COI). I forget how many cells there are in the model. We kept the sizes relatively small ($500 million). As for variable annuities we not only had issue year and issue age distinctions, but we also had asset allocation distinctions. On average, the VA

population had about 70% in equities. We collapsed money market and bonds. We had about $800 million with a 5% roll-up guaranteed minimum death benefit. Then we modeled about $500 million of five-year bank CDs. The universal life (UL) block was backed by callable and non-callable corporate bonds, a variety of different asset qualities, typical insurance company distribution of quality, and mortgage pass-throughs. Just for fun, we put mortgage pass-throughs up against the bank CDs. We're going to look at the risk universe of interest rates. We defined equity risk as the portion of the VA that's associated with domestic and international equity accounts. The portion of the VA that's associated with bond funds or money market is correlated with interest rate risk, credit, mortality, and lapse.

Measuring risk. You can select any number of metrics. It's really a question of what you're comfortable with. We are using earnings at risk. All earnings at risk (EaR) are fairly simply stated. It is undiscounted statutory book profits, so it doesn't include the release of capital. What we're going to do is look at the mean value less the fifth percentile (Chart 1). We're going to look at earnings at risk over one, five, and 10 years. Anything that's stochastic is going through 300 scenarios, and across the entire life of the population, which is 30 years. That is our EaR. We have a risk profile, something that we use a lot. It takes whatever metric you're using, and if you've run it across a number of scenarios, it takes four different pathwise results and rank-orders them from highest to lowest. You have a mean value in here, and you have a fifth percentile value. It actually looks more like a 10th percentile value. I changed the label, but not the graph. You'll have to visualize the fifth percentile. That difference between the mean and the fifth percentile is your amount at risk.

Mean and median are not the same. By the way, the risk metric could be anything you want it to be. It could be our five-year plan, which could be out somewhere in the distribution. It could be the median. All you're trying to do is convey a message about exposure to risk. By using this methodology, you might not look at the fifth percentile. There might be other break points you want to look at and measure. It's one of the nice things about using earnings.

Let's look at the mortality column in Table 1. When we ran the model to capture the mortality risk, all elements were kept constant. That means that we ran a level interest rate scenario, crediting strategy, competitor rate, and lapse reinvestment. The only pathwise distinction was the stochastic process around mortality. Mortality basically had a mean that was equal to whatever mortality table we're assuming, and then we assumed the variance of 5%. The variance was 5% around the qs in the mortality table. At each point or node along the path we're basically flipping coins and deciding whether mortality is some number less than one times the mortality rate, or some number greater than one. Now, what we didn't do in this application, but we have done is overlay trends in mortality improvement. You can build functions around catastrophes. What if we have a reemergence of a bubonic plague? If you want to measure it, you can define it statistically by factoring it into the process. More importantly, you can build in some views and some probabilistic statements about the future for mortality and trends and capture further significant advances in dealing with causes of death like cancer and heart disease. You'll factor that in and have a lot of fun with it. Of course, once you start

doing that, you're going to have to start interpreting the results, which could get somewhat difficult. In this application, we just looked at random fluctuations in the mean mortality table.

TABLE 1
UL RESULTS—5 YEAR EaR

Percentile	Total	Lapse	Credit	Mortality	Interest
1st	$ 0.1	$17.6	$12.9	$5.2	$4.7
5th	3.4	17.8	14.1	7.1	9.0
25th	9.7	18.2	15.3	14.0	15.1
50th	16.4	18.5	16.4	18.4	18.1
75th	21.4	18.7	17.0	23.4	20.0
100th	30.2	19.4	18.1	34.5	23.4
Mean	15.5	18.5	15.9	18.9	16.0
EaR	12.1	0.7	3.0	13.7	11.3

The mean value is 18.9, and this is just a pure sum of book profits over five years. By dividing 18.9 by five, that can be the average annual earnings. The 50th percentile is 18.4. They're fairly close. If you think about it, you would expect them to be close. There's symmetry to the way we're modeling it here, so you wouldn't expect the mean in the 50th percentile to diverge at all. You get a fairly wide range of results, going from $34.5 million to $5.2 million. That amount of risk is 5.2, which is the difference between 18.9 and the EaR 13.7. You can't get hung up on the numbers and the relative values. We just threw stuff in here. When we actually do this work, we spend a lot more time carefully considering the assumptions and the dynamics. Our amount at risk, just purely for mortality, is about 5.2 relative to 18.9. That says that there's a 95% chance that we're going to have base earnings of about 14 or more, but there's a good chance we could lose $8 million. There is a 5% chance we could have a $7.8 million contribution for mortality risk.

Let's move on to the interest rate risk. Mortality is now constant, and we have no default in this model. We should have had default, but we didn't. Interest rates, the crediting strategy, the competitor rates, lapse, and reinvestment are purely stochastic at this point. This is a traditional application on a stand-alone basis that you would be doing in a normal risk-management process. Our mean is at 16.0, and our 50th percentile is 18.1. Some have minimum guarantees and dynamic lapses that are not linear in relationship to changes and between competitor and competitor rates. Some have crediting rates that don't necessarily keep pace with what's going on in terms of your portfolio. Those who have done this work recognize that you're going to get a non-symmetrical distribution, and your mean is going to be less than your standard deviation. I think it will be less than your 50th percentile. When everything was constant, 21.6 ended up being greater than the 18.1. I think that's a typical result. If you ever run your cash-flow testing scenarios, for example, what's the best scenario? The level scenario. In the mean, the 50th percentile is sort of like the level scenario in that context. We have earnings at risk of $11.3 million or 16.0 less 4.7.

Now let's talk about credit risk. Now everything else is constant except the credit level scenario. All the other elements are pretty much set at their default positions, except now default is a stochastic process. We took historical default events by rating category, and fitted distributions to them. You're going to get distributions that are basically consistent with what we've seen in history. That is we have some kind of low level of defaults, and then all of a sudden, we have this big spike up. It is this credit event that's correlated in some way or another with the economic conditions that exist at the time. We're doing our modeling here. We shut off all of the correlations with economic conditions. You can argue that you could correlate credit events with a variety of different indicators, including inflation, GNP, and other things, like the level of interest rates and credit spreads. We didn't want to confuse it. This is purely random, but based on our skewed distributions.

Our mean is 15.9, and our 50th percentile is 16.4. My guess is, if we ran enough scenarios, we would ultimately converge at the same point. Our amount EaR is only $3.0 million. We use millions as our unit. That produces a pretty interesting result. Already we're starting to form an opinion. Interest rate is five times credit, and mortality is four times credit, which is not surprising. Wouldn't you expect mortality and interest rates from a UL product? It also says something about credit. There are a number of ways you could get a specific variation or a specific risk as it relates to lapse. Again, in a level scenario, we sometimes do all the work in an incremental basis, and you can start quantifying your relative contribution to interest rate risk from calls, prepayments, crediting strategies, policyholder behavior, and variation around policyholder behavior. In this case, it's relatively insignificant. All we're doing is bringing in the variation. It was a lot less than we thought it would be, even though we used this methodology.

You add them all up. The total column does not show them added up. It is running the model with everything moving, all at the same time. Everything is stochastic. You get some interesting results. Your mean is 15.5. Your 50th percentile is 16.4. The difference between the 15.5 and the 18.9, is the 18.9 didn't have credit in there. The credit is worth about $3 million, and it didn't have interest rate risk. That's worth another $3 million, so the combined effect of credit and interest rate risks brings you down to around $14 million or so, plus there's the interplay of a lot of different elements.

If you add these up, you would recognize that they're almost twice what you get when you run the entire model, in aggregate.

Let's focus on the UL piece. All I'm doing is recapturing the relative contribution of risk, to come up with a correlated total of $21 million. And the benefit from not having fully integrative, fully correlative, or positively correlated risk elements is a correlation effect of almost $9 million. The correlated total is $12.1 million. We can start ultimately getting the relative contribution to the $15.5 million using some statistical techniques, to decompose the $15.5 million into the contribution that's coming from the various risk elements. We went through a similar process with the VA piece, and what is not surprising is the equity piece is a big component. Interest rate is a smaller component, and lapse and mortality are relatively small,

given the methodology. The uncorrelated total is 15.2, and the correlation effect is a negative 2.6, which is about 12.6. The CD is all interest rate risk.

Then you end up with an uncorrelated total. As for the risk elements, there is $42 million in total, across all the product lines. There is $31 on a correlated basis with a correlation effect of almost $12 million. If, on an uncorrelated basis, you start combining product lines, you get some correlation effects. What's going on in the bond funds is low interest rates are good for a bond fund in a VA. They're not necessarily good for a UL product, so you get some off-saving effects, it depends on whether this product line is mismatched long and whether this product line is mismatched short. But the correlation effect carries over. There is a bit of lapse, again, primarily between the VA and the UL piece. Mortality is uncorrelated in the way we've done it. You end up with an uncorrelated total of $40.2 million, and some of it is coming from the interplay between the risk elements, and some of it is coming from the interplay of the risk elements across the products. You end up with only $17 million, in terms of a correlated total on risk, and you have an idea of where the relative contributions come from. Not surprisingly, interest rate and equity are the two biggest components for the mix of business. Credit is relatively small. If I extrapolate from this, I could start arguing, maybe we should take less interest rate, less equity risk exposure, more credit, and more mortality rate. That's the way you would use this tool, because it allows you to start making some very definitive statements about how much of what type of risk you have, and decide whether or not you're comfortable with it.

Let's discuss the five-year results. We went through the same process, just for one year out, and the same process for 10 years out. And I resisted the urge to annualize these because you really can't. You have to go through some statistical techniques to put them on an apples-to-apples basis. If you want to divide 16.1 by five, and three-point-something, you can. Or, you divide 38.7 by ten and get 3.9. You realize, with the passage of time, the interest rate risk is expanding, and you would naturally expect that. In terms of the expansion of the scenarios, as well as the fact that any surrender charges are gone, this is just an in-force block. On the equity exposure, you get the benefits of some time diversification. You certainly get the benefits of time diversification. A one-year credit event is a big deal. Over 10 years, how many big credit events are you going to have? It certainly is not historically 10. The way we've looked at credit risk, you're going to get the benefits of some time diversification here. I forgot to mention, when you do credit, you could start factoring. We're modeling every bond seriatim and we're going through a stochastic process in terms of deciding whether or not that bond defaults or doesn't default. But you can also segment the bond population in terms of seasonality. You could have different credit distributions based on whether or not the bond is a seasoned bond, how far from issue it has been, and so on. There's some data that can help you to get there. Lapse becomes a bigger issue over time, and mortality stays fairly stable. And if you start doing the math, you would see that, on an annualized basis, it kind of offsets each other. Some are getting worse with time and others are getting better. You end up with about three and change a year, if you annualize the results.

I believe that the marketplace realities mandate that we move to this new risk management paradigm, beyond interest rate risk. Management has a clear need for the information. I have spent more time with chief financial officers and CEOs, who constantly complain about the absence of clearly understandable information about the company's overall risk exposure. Are they so frustrated that they come down from wherever they sit and somehow mandate it? No. But there's a frustration that the information isn't there. Finally, it's achievable, and it's just tremendously powerful.

Ms. Bennett: Even though Frank and I did not have a chance to speak to each other about our presentations, you're going to see a tremendous amount of similarity in our comments as they relate to risk management.

I agree with Frank that we need to start emphasizing risk management within our organizations. It is important to think in terms of enterprise-wide or holistic risk management. I am going to talk about a method for managing interest rate risk. While I agree with Frank that eventually we need to move risk management beyond interest rate risk, I think we still have some work to do in the area of interest rate risk management. I am going to take a step back and focus on a method for managing the interest rate risk. In my former company, the interest rate risk was the single largest risk facing the company, given the mix of business.

Risk management is getting a lot more attention nowadays. There are more conferences specifically centered around risk management. Dedicated risk management positions are being created in more insurance companies in the investment area or the actuarial departments. With all this attention, it almost seems as if risk management is becoming a new discipline, or something that we've never thought about before in managing insurance companies. However, risk management is not anything new for the actuary. We've been doing claims expense and interest rate risk management for some time. I think the difference is that actuaries have taken a functional view of risk rather than an enterprise-wide view of risk. We've studied mortality and modified the claims or the underwriting practices. We've studied defaults and changed the underwriting for defaults. We've taken a functional approach, but now we're seeing the need for risk management, which encompasses the entire organization. Risk management must be dynamic and encompass all the risks. In accounting terms, we have to look at both sides of the balance sheet.

While we've been doing cash-flow testing and asset/liability management (ALM) for 15 or 20 years, we still recognize the need to do better in financial management. We are seeing a growing acceptance of the need for integrated ALM due to the increased sophistication of products and consumers, and the different asset types available. Integrated ALM, or financial management, is becoming a necessity in managing a company's financial position. While ALM has tended to be the domain of larger companies, even the smaller companies are starting to recognize that they can't afford to ignore the risks either.

When we think about risk management, it's not that the actuaries have been just standing still and not doing anything. The actuaries, too, have become more

sophisticated in the area of financial management. Actuaries are more often being appealed to not only add up the financial results, but to also help navigate the company toward achieving the desired financial results.

We have a lot of exciting prospects on the horizon with risk management and financial management, but we have to ask ourselves whether or not we have the tools and the structures in place to add value to the risk management process. Are the science and the concepts of risk management developed adequately enough to be applied to the life insurance industry? Do we have the technology? Do we have comprehensive corporate models that are sufficiently robust to be used responsibly in the risk management process? Finally, do we have the software and the machine power to run the modeling applications?

Before we answer these questions, we have to ask ourselves what we're trying to accomplish here. You may have noticed that I tend to use the words *ALM, integrated financial management,* and *interest rate risk management* interchangeably. Essentially, I think these terms represent the processes of understanding the financial risk dimensions of the organization. We're making decisions to manage the company's capital, earnings, or risk position. Capital represents the funds needed to cover the risks assumed by the organization. The earnings represent the change in capital from period to period, and risk represents the variability in income from period to period. We make risk management decisions every day. Just take a moment and think about some of the things that you do every day. Some of you might be involved with setting crediting rates or pricing products. I'm setting the crediting rates. The basic question you're asking is whether or not you're trying to increase or maintain market share. If you're trying to increase market share, you will probably have to squeeze earnings. If you're trying to maintain market share, then you're probably satisfied with your current financial position. Perhaps, in the past, you've been aggressively setting past crediting rates in an attempt to increase market share. As such, you might need to increase margins to contribute to surplus. Regardless of whether you are setting crediting rates and pricing products, you are making decisions about capital, earnings, and risk. If you are involved with setting investment policy (and some of you may be considering asset types with both fixed income and equity components), you are making decisions about the trade-off between short-term income and long-term capital appreciation. You're making some bets about how you think those assets are going to perform in different interest environments.

You're making decisions regarding assets to purchase in the event of liquidity risk materializing. You might be wondering whether or not your company should be purchasing interest-rate floors, given what's happening with interest rates. Perhaps some of you are involved with the business-planning process and wondering about the future of fixed-account products. Some think fixed-account products are a thing of the past, and the future lies solely with variable products. In strategic planning, you're making some decisions about changes to your product and customer base. And finally, some of you may be involved in setting financial management objectives. There are some common measures used by many companies such as GAAP return on equity, risk-based capital ratios, and growth targets. Are these objectives well-defined, consistently defined, and achievable?

My point in posing these questions to you is that your company is managing the financial risk. Risk management may be a conscious or unconscious process; a haphazard part of management. If we want to better understand financial management, we have to ask if we need to build a better mousetrap. What are we looking for in terms of financial management? As a corporate actuary, I was often faced with many of these questions that I posed earlier, and too often, I had to admit that I really couldn't answer the questions. I understood the nature of the question, but I didn't feel confident to direct certain actions because I didn't have sufficient information. I needed to strengthen the financial management infrastructure so that better financial risk management decisions could be made. I needed a financial management process that would allow me to manage the interest rate risk. Now, in order to manage the interest rate risk, I had to first be able to measure the risk. I needed a financial management process that would allow me to evaluate alternative investment and capital strategies. I needed a financial management process that would allow me to refine profit expectations by line of business and for the total company. And finally, I needed a financial management process that would allow me to evaluate performance of the asset, product, and corporate managers.

A big obstacle in leveraging the existing ALM systems was that the current financial reporting techniques for allocating investment results to the product lines commingled the contributions of the asset, product, and corporate managers. We could all discuss the numerous limitations of statutory and GAAP accounting, but right now, statutory and GAAP accounting define our operating environment. Many of our management information systems are based on meeting statutory or GAAP requirements. Many of our resources, whether they be systems or human resources, are dedicated to preparing and explaining GAAP and statutory accounting results, often leaving us insufficient time to analyze the underlying financial economics.

It became clear to me that the process for allocating investment results to the product lines had to be reconfigured to facilitate the management of the interest rate risk, and ultimately, the development of a more practical ALM system. There are a lot of different ideas for building an ALM (ALM) system. Our profession has been grappling with several of these concepts, such as value at risk, distributable earnings, and others. All companies want to capitalize on the investment that they have made in cash-flow testing systems and to use the information from these systems to manage the company.

Towards that end, I'd like to offer my contribution to this whole discussion. I have spent the last 18 months developing a transfer pricing paradigm for managing the financial condition of a life insurance company. Now, the term transfer pricing is starting to get some recognition. It's a term or a concept borrowed from the banking industry. Let me define transfer pricing. Essentially, transfer pricing is the inter-company reinsurance of the interest rate risk. In transfer pricing, the product managers cede or transfer the interest rate risk to the corporate line. We'll discuss transfer pricing in detail, but you might be thinking that transfer pricing is merely an accounting approach. I can assure you that transfer pricing is a lot more than

an accounting system. Once you have a transfer pricing paradigm in place, many of the obstacles that impede economic analysis are removed.

Let's consider the basic problem that we're faced with as an insurance company. We issue products with an expected interest rate risk profile. In a typical price behavior curve, price is measured as the present value of cash flows. The question we're faced with is, which assets and how much capital should back these product liabilities? Most companies have some type of segmented asset allocation approach for allocating assets and capital to their liabilities. Either portions of assets or entire classes of assets back certain types of products. You're faced with an ALM paradigm with segmented asset allocation that has assets moving from 1,800 to 1,400 and liabilities moving from 1,350 to about 1,200 in the price behavior curve.

We need to evaluate this ALM paradigm. Is the investment strategy appropriate, given the product risk? Is there an acceptable level of capital backing this product? I have to admit, I've been looking at these kinds of graphs for several years, and I'm not sure how to answer these questions. The question to pose is whether we're doing the best thing. Is this an optimal investment strategy or an optimal capital allocation strategy?

Now, with transfer pricing, we redesign the ALM paradigm. Within transfer pricing, a centralized corporate risk function is established. The interest rate risk is transferred from the product lines to the corporate line. The corporate line will contain the assets that are not necessary to support the product obligations. The line of business assets will be equal to the reserves plus any target capital. For the purpose of this demonstration, I define target capital according to the standard NAIC formula. The C-3 risk capital for interest rate risk is allocated to the corporate line. The corporate product line contains all of the C-3 risk plus any free capital for new ventures.

With transfer pricing, the first thing to do is change the dollar amount of assets that support the product lines. Next, the composition of the assets is changed. We're changing the composition of those assets by creating synthetic asset portfolios. With the dollar amount and composition of the assets reallocated, the income statement entries change accordingly. The investment income, capital gains and all the associated items are now allocated to be consistent with the responsibilities of the product, asset, and corporate managers.

Within transfer pricing, the construction of the synthetic asset portfolios is an important step. With transfer pricing, many of the suboptimal results of a segmented asset allocation are avoided. The synthetic asset portfolios match the interest rate risk profile of the liabilities, as represented by the price behavior curves. The synthetic asset portfolios contain non-callable bonds of various maturities, and interest rate derivatives such as caps and floors. The synthetic asset portfolios are constructed from a linear programming model.

Within the transfer pricing paradigm we have reconstructed the ALM paradigm. We're still working with the same liability curve but now, we have a synthetic asset portfolio that matches the product liabilities. The assets and liabilities are matched

along the entire price curve, which is well beyond duration or convexity matching. This ALM paradigm insulates the product manager from the interest rate risk. The synthetic asset portfolios will now move in tandem with the liabilities. The product managers are no longer influenced, either positively or negatively, by the actions of the asset manager. The product manager will be able to concentrate on managing the product. We're in a position to evaluate what the product manager has done in his or her own right, independent of the actions of others. It's going to be easier to evaluate the asset managers' actions and have a basis for setting the investment strategy for the entire general account asset portfolio.

Transfer pricing can eliminate some of the finger-pointing that can go on within an organization. In expanding earnings, a product manager could cite lower investment income due to increased prepayment activity or non-income-producing equity investments as reason for falling short of a goal. The asset people can speculate that the products are mispriced and create an incentive for policyholders to exercise product options. While some of those comments are legitimate, the problem is that it's hard to figure out what to do with the products or investments to turn things around. With transfer pricing you can break through some of the finger-pointing that can take place. With transfer pricing, the product and asset managers can be held accountable for their own actions and not blame other people.

As I said earlier, transfer pricing, at first blush, looks like an asset or capital allocation process, but it's really a whole lot more. Once you have a transfer pricing system set up, the interest rate risk is managed in the corporate line for the benefit of the enterprise. There are many applications for managing the financial position of your organization. The first application, and maybe the most important, is that the interest rate risk can be quantified. As product options are granted and exercised, the cost of these options shows up in the income statement. As asset classes with options or equity components are purchased, the risk and return of those asset classes can be measured.

The investment strategy can be established with direct recognition of the product risk and free capital for the organization. The synthetic portfolios represent a bifurcation of the actual assets purchased into the fixed and options components. We're not creating "fake" assets but we're bifurcating the actual assets that have been purchased into the fixed piece for more predictable product line cash flow. We moved the options piece to the corporate line to centrally manage the interest rate risk.

In setting crediting rates and pricing products, we can now directly recognize the interest rate risk. Not every company uses ALM and stochastic techniques in the pricing of their products. Many times we hope certain risks offset each other, or maybe we're just not sure, but now we can directly recognize the interest rate risk in setting the crediting rates. We're in a position now to evaluate the risk and return of alternative strategies, whether they are product, asset, or capital. We can evaluate the contributions of the asset, product, and corporate managers.

With transfer pricing, we have the basis for establishing customized benchmarks for asset performance. These synthetic asset portfolios provide a minimum threshold for performance evaluation. At a minimum we have to purchase assets with cash flow sufficient to cover the product obligations. We can go to our asset manager and illustrate how we could buy this pool of non callable bonds and caps and floors, and this asset pool will provide for the promises made to our policyholders. Now, I still expect the asset manager to invest in assets with better returns, but at a minimum, we know that they have to exceed the synthetic threshold.

In this session we're focusing on risk management. I want to emphasize that many risks can be evaluated using transfer pricing. We can look at an investment strategy that might involve taking on more credit risk, and we can evaluate the risk and return trade-offs of having purchased bonds of lower investment quality. You can evaluate the credit risk in terms of the trade-off between current income and capital loss. You can evaluate equity risks. With many investment products, you're betting on the equity markets. The investments represent a trade-off of current income in exchange for future appreciation. You're taking a bet on current earnings, and you're also taking a bet on how the value of your company is going to change. With transfer pricing, you're going to be in a position to quantify the magnitude of the bet.

On the product side, you're going to be able to evaluate the costs of certain options and guarantees embedded in the contract. There's a lot of pressure to add many bells and whistles to new products. We're often not really sure how to price the options. We can quantify the cost of these options and guarantees on short-term versus long-term earnings.

The transfer pricing paradigm in effect gives us an asset/liability model that allows us to ask a lot of what-if questions, and answer the questions in terms of the trade-offs between short-term earnings and capital appreciation. I think the transfer pricing framework gives us new tools to manage the financial condition of the firm. With transfer pricing, we will have more predictable line-of-business income statements, where the contributions of the asset and product managers are disaggregated. We can project the company's financial performance based on a range of expected economic conditions, and measure the projected risk-and-return trade-offs of alternative strategies. With transfer pricing, we'll have a basis for establishing product pricing expectations, not only for new business, but also for overall corporate company profitability which directly reflects the risks of the cost of capital. As I mentioned earlier, I have spent the last 18 months implementing and designing the transfer pricing approach at Minnesota Life. I am happy to report, and frankly a little bit relieved, but this theory that I've laid out here actually works in practice. This isn't just theory; this really does work! There are many different uses of a transfer pricing paradigm, so it's important, once you get the system in place, to concentrate on the most critical risks for your organization. You might need more sophisticated analysis of product strategy or pricing. You might need more sophisticated analysis with investment strategy. If you invest the time in setting up an overall risk management process, which involves transfer pricing, you'll have a very powerful tool within your organization. I've become very, very convinced that the transfer pricing paradigm is the elusive answer that the

insurance industry has been seeking to manage the interest rate risk and move us well beyond regulatory cash-flow testing.

From The Floor: I have a question for Frank Sabatini regarding the case study. When I look at the VAR number or EAR number, I'm interested in seeing what my exposure is, given my current ALM strategy or hedging strategy. What is done with hedging on the asset side?

Mr. Sabatini: Future hedging has been reflected. The numbers represent a straight application of a particular asset inventory at the start of the projection process, and then a reinvestment strategy. There is nothing dynamic, although the facility that we're using could lend itself to build in dynamic strategies.

Mr. John S. Bath: I'm interested in the synthetic asset portfolio approach. Do you set some spread over Treasuries that varies by product depending upon liquidity issues?

Ms. Bennett: The question relates to construction of the synthetic portfolios by product type. We set up six different major liability segments, essentially setting up six small companies. The synthetic assets are different and completely independent for the deferred annuity segment versus the individual live segment, et cetera. We set up a fairly complex linear programming model with historical prices for the synthetic assets. The approach we took was more on the balance sheet side, rather than on the income statement side. A difference between transfer pricing in the banks and the insurance companies is that banking transfer pricing is defined on the income statement side, where you define the risk or the spread for having transferred the risk, creating a cost to the product lines. I don't explicitly define a risk charge; the risk charge is implicitly defined in the construction of the synthetic asset portfolios.

Mr. Peter D. Tilley: Frank, in the comprehensive system that you have where you're looking at the overall total, stochastically modeling interest rate changes, equity markets, mortality fluctuations, how did you come up with the number 300 for scenarios? It just seems to me that there's so much stochastic things going on that you'd need more.

Mr. Sabatini: In a production mode, there is more. We were doing a number of different iterations, just to get ready for the presentation, and 300 was the number we chose because for quick turnaround. In a production context, you really need to figure out how many scenarios you have to run to get convergence.

CHART 1
EARNINGS AT RISK (EaR)