RECORD, Volume 29, No. 3*

Orlando Annual Meeting
October 26–29, 2003

Session 65TS
"One Shoe Doesn't Fit All": Using the Appropriate Pricing Measure

Track: Product Development, Education & Research

Moderator: NANCY WESTFALL WININGS
Instructors: ANDREW D. CHUA
NANCY WESTFALL WININGS

Summary: There are many measures of profitability in the life insurance and annuity product development processes, including internal rate of return (IRR), return on equity (ROE), return on assets (ROA), profit margin and break-even year. Which are the most appropriate, and what are they telling you? How does the answer differ by product, and why? How might profit targets differ under changing circumstances?

MS. NANCY WESTFALL WININGS: I work for Milliman USA. Our topic is going to be profit measures that are appropriate to use for various products, and their pitfalls.

My co-speaker is Andrew Chua, who is in product development and management at Legacy Marketing Group. His focus is on annuities. He previously was the senior vice president of risk management at ING Direct. I'll be discussing the life side. We're going to play off of each other a little bit, because this topic is something that pervades both sides. We'll both have things to say about the different profit measures that we have.

We will not be discussing which profit measures are *the* ones to use or what levels you should be using, but rather what the measures are and the reasons to pick one

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

over another for a particular product. We have a couple of goals that we would like to accomplish during this presentation. First of all, we would like to identify and describe the common profit measures that are used in the life insurance industry. We're going to focus on the ones that are primarily used, but I know that there are other ones out there that you may want to discuss. We're going to focus particularly on the issues regarding profit measures and compare and contrast information that each profit measure provides. We're going to try to explain why you might pick a particular profit measure for a certain product. The main profit measures that we will focus on are IRR, ROE, embedded value, ROA, profit margin, and risk-adjusted return on capital (RAROC).

We also will touch on break-even analysis, surplus strain and different types of sensitivity analysis—whether deterministic or stochastic. We will explain what companies might do to try to understand the sensitivity of the particular profit measure that they have chosen.

MR. ANDREW D. CHUA: IRR, in a nutshell, is the interest rate that equates the investment that's put into the product against the payout from the product. Investment, typically, is product strain and target surplus. Product strain is the fund that you receive, adding interest to it, paying out taxes, paying out benefits and expenses, and comparing that to your reserve. The difference is the strain.

Target surplus, typically, is prescribed by management. The decision is driven by either economic capital, regulatory requirements or a rating agency. Regulatory considerations involve risk-based capital (RBC). For some of you who might not know this, you'll hear the terms "authorized level" and "RBC factor." The authorized level is 50 percent of the RBC factor. Sometimes there is confusion there. You'll hear that it is 250 percent of RBC, and people will scratch their heads when somebody throws out a number like 125 percent of RBC factor. And the reason is that kind of translation. The rating agencies are AM Best Company, Standard and Poor's (S&P) or Moody's Investors Service. S&P has a description of its capital-adequacy model, and you can get a copy from their Web site. AM Best has what is called a financial strength rating (FSR), and again, you can get some more information about that from the Web site.

You can compare an FSR rating to an S&P rating. I think that you will find it very valuable when talking to management, as you translate from one to another. For example, you'll find that an AM Best "A" is essentially the same as an S&P "A." And an AM Best "A+" is an "AA" with S&P. Another table that you might find useful for purposes of understanding capital requirement is in a book by Farbozzi. He summarized all of the ratings from the various rating agencies into one table. You'll find that, for comparing one rating agency to another, it's a very useful piece to have.

If you look at IRR, it's similar to investing in bonds—except that with insurance product, there is more risk. There's more investment. The cash flow has more

optionality. Because of all of the risk associated with optionality as well as the investment, one would expect the rate of return on life insurance products to be higher than investment-grade bonds.

The preferred situation when determining IRR is to see a negative investment in the first year and then a bunch of payouts in the future. With all of the nuances that now exist in products, that's very difficult to accomplish. What happens is that you'll see a mix of positive and negative amounts. It's not only a mathematical problem, but conceptually it is a problem in measuring profitability. The level of certainty attached to the cash flow owed to a policyholder is very high. It's at least equal to the rating of the insurance company. If the rating is AA, then that's probably the level of certainty associated with that cash flow. The discount rate used should be consistent with that level of certainty.

Because of optionality—the uncertainty, the fact that we're using assumptions that go pretty far into the future—we have to adjust for the risk embedded into the cash flow that the insurance company receives. Hence the discount rate used should be higher than the one used for discounting cash flows owed to policyholders. David Becker wrote a paper about this, and that is what his approach is trying to communicate. The two cash flows should not be discounted in the same way.

The first time you go through the design of a product, you may find this mix of positive and negative cash flows. You either can change the benefits, change the assumptions appropriately or change the reserves and pre-fund future losses. These are alternatives to the Becker approach.

Pros and cons—we'll start off with the "pros" side. IRR—if you properly adjust the investment into the product consistent with the risk—allows for good comparison among products. Because of the nature of the calculation, it's very similar to an investment in the bond. It helps ensure that the returns are adequate, especially if you properly adjust the investment. It's easy to understand. It's like investing in a bond. The methodology itself allows you to make adjustment for risk. The challenge is making those adjustments. That's where the complexity comes in. For example, especially with annuities, they could adjust to reflect the cost of reinsurance or reflect the charge associated with investing in the capital market or hedging against the capital market.

IRR by itself is very useful. But it's incomplete. It doesn't tell you whether you're investing one dollar or a million dollars. It doesn't tell you whether the positive payout from the product is going to occur starting in year 10 or in the fourth year, so it has to be supplemented by additional information.

The typical approach that I've seen is the tendency to follow the target surplus requirement prescribed by RBC, AM Best, or one of the other rating agencies, without really looking at the economic capital associated with the product. And if we don't do that, then the IRR calculation may be inappropriate. The best example

would be for guaranteed minimum accumulation benefits (GMABs) in variable annuities. Until fairly recently, there was no capital required for it. Yet we know that it's a very risky product feature.

MS. WININGS: I am going to go over some examples of how to calculate IRR with the intention of establishing the fact that there are some sensitivities in calculating IRRs. Your product design can make it such that the IRR is inappropriate or you need to make adjustments. So if you are reporting to upper management, you should be able to explain what's going on with the numbers. Or if you are upper management, you can ask the right questions in order to understand what is going on.

In Table 1, I show the results of two different products that I designed. The first one has a relatively high reserve requirement, and that has subsequent positive distributable profits. The second example is identical for years two and later, but has a lower reserve requirement. As would be expected, the IRR is slightly higher. The point being that if we would eliminate the first year of loss, we would have an infinite IRR that would be somewhat meaningless. So for those products that might have very low reserves, perhaps the IRR is not the only measure that you should be looking at. There should be alternatives; maybe this isn't the right measure. This figure shows a basic example. Essentially, the way that we calculated this was to discount at a constant rate all the future cash flows back to time zero.

Table 1

IRR	**7.30%**	**10.05%**
	Example 1	Example 2
	Distributable	Distributable
Duration	Profits	Profits
1	(1,265.80)	(949.35)
2	146.76	146.76
3	103.37	103.37
•	•	•
•	•	•
14	166.33	166.33
15	177.58	177.58
16	236.39	236.39
17	303.71	303.71
18	332.82	332.82
19	366.29	366.29
20	403.85	403.85
21	306.40	306.40
•	•	•
•	•	•
29	0.45	0.45
30	0.15	0.15

Only one change in sign throughout stream

In Table 2, I tried to demonstrate the generalized IRR approach that was written about in a paper by David Becker. The paper goes through several lengthy examples. And in this particular one, we are taking a look at the accumulation of cash flows. In this example, I have a very large first-year loss, but I have subsequent losses in later years. Because there are multiple sign changes in this stream, I have the potential for multiple roots. The concern is that I am not properly recognizing the fact that I have these later losses. If you do have later losses, you need to recognize the fact that you need to set aside funds to handle that. The way this particular example works is that we have to say that at some point, before we hit our losses, we have to start accruing money into an account. When those losses hit, we will have enough money.

Table 2

IRR **10.18%**

Generalized IRR **10.04%**

Duration	Distributable Profit	O/S Balance	Profits Released
1	(1,305.81)	(1,305.81)	(1,305.81)
2	283.27	(1,153.70)	283.27
3	233.58	(1,036.00)	233.58
•	•	•	•
•	•	•	•
14	194.17	(74.10)	194.17
15	171.95	90.41	81.54
16	155.67	251.50	-
17	88.78	355.38	-
18	(77.73)	298.97	-
19	(202.46)	114.45	-
20	(201.06)	(79.75)	(79.75)
21	69.13	(18.62)	69.13
•	•	•	•
•	•	•	•
29	0.83	(0.25)	0.83
30	0.27	(0.00)	0.27

Accumulation Rate for Negative Balance = 6%

What ends up happening is that you are solving for that discount rate to those positive cash flows. So that when you discount the pieces at 6 percent, to fund your future losses, you end up at the zero balance. The Becker method is trying to deal with negative cash flows in future years. There is another approach to it. Instead of looking at this as an accumulated basis (which is what I've done), they discount back to time zero. But we get to the same spot.

In Table 3, I'm trying to demonstrate that you really need to understand how your product works. If you were to look blindly at the most-right-hand column of this example, you may be thrilled that you had a product that is delivering a 28.7 percent IRR. And in fact, the situation is that you have a lot of very large negative later-year losses. You're relying on the fact that you're going to be setting aside some money, and it's earning 6 percent. It gives you a skewed answer. And so, it can be sensitive to the rate that you are going to be earning; the positive cash flows offset the future negatives. You may want to set up a different set of reserves or tweak your product design. There can be a lot of variability between the standard IRR and the Becker IRR.

Table 3

IRR 2.37%

Generalized IRR 28.74%

Duration	Distributable Profit	O/S Balance	Profits Released
1	(387.83)	(387.83)	(387.83)
2	188.27	(311.06)	188.27
3	168.48	(232.00)	168.48
•	•	•	•
•	•	•	•
14	74.85	1,195.96	-
15	42.77	1,310.49	-
16	(23.16)	1,365.95	-
17	(144.51)	1,303.40	-
18	(334.92)	1,046.68	-
19	(487.00)	622.48	-
20	(516.21)	143.61	-
21	(170.82)	(18.59)	(18.59)
•	•	•	•
•	•	•	•
29	1.74	(0.70)	1.74
30	0.79	(0.12)	0.79

Accumulation Rate for Negative Balance = 6%

I ran across the example in Chart 1 in some work that I was doing on a return-of-premium term-insurance product. For those of you who aren't familiar with that product, this could be considered a term product. And over time, you get a return of your premium. So let's say that at year 30, you would get 100 percent of the premium paid if you were to persist to that point in time. Prior to that point, you would get something less than 100 percent of your accumulated premiums. And right now, a lot of the products are designed so that it's 0 percent up to year five. And then it gradually increases between years five and 10, a little more between years 10 and 15, and so forth. But the point is that I had a 20-year product, and in the year 15 there was return of premium on it. So I had these high reserves that kicked in because I had to hold at least the cash value. I had very high reserves that kicked in, but I didn't have very much variability in those first two years. I just changed one or two assumptions, and I got very similar year-by-year profit patterns.

Table 4 shows the results. Product A is the one that starts with a higher first-year loss in its absolute magnitude and slightly lower positives in the later years. Product B is the one that has a smaller first-year loss, and slightly higher positives. It doesn't look like there's much difference between these. At first, you would say that you should have similar IRRs. And looking at the generalized IRR for these two different products, you see that there's a base difference in the generalized IRR. In later years, you can see that they are identical. The only difference is related to the first-year loss and then the subsequent positives. That's the only reason that they could have these differences in the IRR.

Table 4

Accumulation Rate for Negative Flows	Generalized IRR	
	A	B
4%	-8.9%	11.7%
6%	.5%	17.3%
8%	6.3%	20.8%

I think that it's important to realize that you can have a wide variability due to the fact that you have a different first-year loss and slightly different positive gains. But it's a pretty wide difference. And if you had looked only at the one versus the other, without tweaking one assumption and taking a look at the sensitivity of your results to a variance in one of your assumptions, you may come to the conclusion that this is a great product. But Column A shows a very different picture of how this product might work. It is very sensitive to the assumptions. Even taking a look at whether or not you put your funds into a 4, 6 or 8 percent account, you still end up with similar conclusions. One gives you much higher returns than the other.

MR. CHUA: The SOA Risk Management Task Force sponsored a survey in late 2002. The survey was intended to look at risk measures broken down by several different criteria:

1. Area of practice—whether you're in valuation, risk management, pricing, marketing, and investment; and how popular those varying risk-management issues are to those various individuals.
2. Type of insurance organization—insurance, consulting, reinsurance and others.
3. Product type—whole-life, term, variable life, universal life, fixed, and variable annuities.

They asked, within that breakdown, how the various individuals are approaching risk adjustment to their profit measures. And the choices were assumption pads, assumption stress testing, capital allocation, target surplus, risk-adjusted profit target or stochastic analysis. The survey showed that the two primary risk adjustments to IRR are through capital allocation and assumption stress testing. Other significant adjustment methods are through a risk-adjusted target-profit measure, as well as stochastic scenario analysis.

Particularly with respect to life and annuities, the capital allocation and assumption stress testing were just about equal. With respect to individual health, the dominant risk-adjustment methodology is through assumption stress testing by a wide margin.

ROE is simply GAAP income over GAAP equity. We'll try to tie this to IRR. It's typically an annual measure. In order to apply it to insurance products that tend to be medium- to long-term, we have to make certain adjustments. Simple average (sum of income over sum of equity), discounted average (present value of income over present value of equity) and production weighted (layers of new business) seem to be the most popular.

ROE is required by the SEC for stock companies. Experience seems to show that it's a more appropriate measure for inforce blocks of business. Simply because of the nature of new business, the pattern of earnings tends to be somewhat steep, especially in certain products. If you have a large block of business, the ROE calculation from year to year is more stable.

The pros are that it's well known and understood. It is the de facto profit standard for Wall Street, and it measures effectiveness of capital or equity. I want to emphasize capital or equity, as opposed to investment. The GAAP equity is not equal to the investment in the product. Because it measures effectiveness of capital, it does not properly measure the time/value money; it does not measure return relative to the investment. In going from IRR to ROE, there are a number of adjustments. And lastly, it is highly dependent on the product design and the deferred-acquisition cost (DAC) methodology that you use.

Table 5 shows an example of a fixed annuity. It uses an earned rate of 6.5 percent with a credited rate of 4.5 percent. It assumes target surplus at 5 percent, a 35 percent tax rate, and tax reserves and statutory reserves equal to cash values. The adjustment is simply the deferral of acquisition costs. It assumes a total compensation of about 8.25 percent. The liability adjustment is the surrender charge. GAAP causes you to set up reserves equal to account value. The difference is an adjustment to equity. There are other adjustments to GAAP that I have not reflected here. There are two that I think are significant. One is the tax piece of it. On the statutory side, you use tax reserves and you have DAC tax. On the GAAP side, you take your pre-tax income, apply 35 percent to it, and that's the tax. The difference between the two is either an asset or a liability item.

Table 5

	Stat Balance Sheet			Adjustments			GAAP Balance Sheet		
Time	Stat Assets	Stat Rx	Target Surplus	Assets	Liability	Equity	GAAP Assets	GAAP Rx	GAAP Equity
0	976.50	930.00	46.50	82.50	70.00	12.50	1,059.00	1,000.00	59.00
1	990.03	942.89	47.14	73.37	70.97	2.40	1,063.40	1,013.86	49.54
2	1,002.17	954.44	47.72	64.28	60.92	3.36	1,066.45	1,015.37	51.08
3	1,001.84	954.14	47.71	54.71	50.22	4.49	1,056.55	1,004.35	52.20
4	989.97	942.83	47.14	44.91	39.28	5.62	1,034.88	982.12	52.76
5	965.88	919.88	45.99	35.11	28.45	6.66	1,000.99	948.33	52.65
6	930.35	886.04	44.30	25.57	18.08	7.48	955.91	904.13	51.79
7	884.59	842.46	42.12	16.51	8.51	8.00	901.09	850.97	50.12
8	547.35	521.29	26.06	7.41	0.00	7.41	554.76	521.29	33.47
9	473.55	451.00	22.55	3.91	0.00	3.91	477.45	451.00	26.46
10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

The other major adjustment that you might see is in the investment income. Particularly with FASB-115, you'll see a difference in the investment income that shows up on your GAAP versus your stat. If you have a lot of merger acquisitions, you have that additional adjustment.

Table 6 shows the statutory net income. I want to emphasize that what is shown here is statutory net income, not distributable earnings. The difference between the two is the change in target surplus. I did include the interest on target surplus into the statutory net income. Take out whatever expenses you can defer and add in the piece that you have to amortize into income. You make an adjustment for the stat to GAAP reserve difference, and you get GAAP net income and GAAP equity. Year-by-year ROE can now be determined. Average ROEs computed using the simple and discounted methods are also readily determined.

Table 6

Pol Year	Stat Net Income	Dac Cap & Amort	Rx Adjustment	Gaap Net Income	Avg GAAP Equity	ROE
1	4.15	73.37	(70.97)	6.54	54.27	12.06%
2	5.55	(9.09)	10.05	6.51	50.31	12.94%
3	5.73	(9.57)	10.70	6.86	51.64	13.29%
4	5.90	(9.80)	10.93	7.03	52.48	13.39%
5	5.98	(9.80)	10.83	7.02	52.71	13.32%
6	6.01	(9.54)	10.37	6.84	52.22	13.09%
7	5.98	(9.06)	9.57	6.49	50.95	12.74%
8	6.06	(9.10)	8.51	5.47	41.80	13.09%
9	7.25	(3.50)	0.00	3.75	29.96	12.50%
10	6.52	(3.91)	0.00	2.61	13.23	19.77%
Simple:				$59.12	$449.58	13.15%
Discounted:				$43.84	$335.21	13.08%

Please note the high ROE for the last year. I want to emphasize that if you simply apply the typical methodology, which is the average of the beginning and ending GAAP equity, you may come out with a number like this. This example assumes a fixed annuity for 10 years. Everybody lapses at the end of 10 years. If you properly reflect the true average equity that is invested in the product, this number, instead of 19.77 percent, will be somewhere in the neighborhood of 13 percent and change.

From the survey for individual life, the two risk adjustments to ROE that seem to be most popular are capital allocation and assumption stress testing. For individual annuity, it's the same two—except that stochastic scenario analysis is added to the group. For individual health, it's all of the above.

For ROA, we'll talk about when you would use this profit margin and under what conditions. We'll try to link that to what we have just done on the ROE side. ROA is the average spread on assets net of the rate credited to the policyholder, expenses and taxes. It's very popular with spread business especially when there is not much capital invested into the product.

There are several definitions. It can be GAAP or stat, before or after tax, and it can be either distributable, or include target surplus or profit only. Typically, the assets are averaged during the years. It's one-half of the beginning and ending assets. ROA, on a GAAP basis, is ROE times a factor equal to equity divided by asset.

ROA is a well understood profit measure. It's popular for spread-based products. It is the most appropriate profit measure if the investment in the product is fairly low. That investment is typically based on prescribed target surplus as opposed to what

might be appropriate. Take variable annuity, for example. I think that in the past 18 months, we've all seen the impact of revenue volatility to the insurance company, either due to DAC write-down or the fact that their asset level was all over the place. And yet, generally, one does not price for that in the product. The fact is that your income can be very volatile in a variable annuity product, as opposed to a fixed annuity product.

If the reserve is a very small number, ROA is meaningless. With no assets, you don't have a denominator. Another challenge with ROA is the lack of consistent definition. The target spread seems to vary by products. It's because of the fact that it's more subjective, in my opinion, as opposed to a rate of return, for which you can go to the capital market and get a sense of what sort of risk premium is appropriate. In the case of a spread, you have to look at what other products are offering as an appropriate spread.

Table 7 is an ROA example based on the earlier ROE example. All I've done is taken the GAAP net income—both the equity and the assets on the GAAP side—and applied the formula either directly by taking GAAP net income divided by average GAAP asset, or using the adjustment of equity over asset to go from ROE to ROA.

Table 7

Time	GAAP Net Income	Avg GAAP Assets	Avg GAAP Equity	Asset/Equity	* ROA	= ROE
0						
1	6.54	$ 1,061.20	$ 54.27	19.55	0.62%	12.06%
2	6.51	1,064.93	50.31	21.17	0.61%	12.94%
3	6.86	1,061.50	51.64	20.56	0.65%	13.29%
4	7.03	1,045.72	52.48	19.93	0.67%	13.39%
5	7.02	1,017.93	52.71	19.31	0.69%	13.32%
6	6.84	978.45	52.22	18.74	0.70%	13.09%
7	6.49	928.50	50.95	18.22	0.70%	12.74%
8	5.47	727.93	41.80	17.42	0.75%	13.09%
9	3.75	516.11	29.96	17.22	0.73%	12.50%
10	2.61	238.73	13.23	18.05	1.10%	19.77%
Simple:	$59.12	$8,640.99	$ 449.58	19.22	0.68%	13.15%
Discounted:	$43.84	$6,497.81	$ 335.21	19.38	0.67%	13.08%

I want to emphasize that there are typically a lot of different adjustments. This is a simplistic illustration. With regard to the survey, the ROA results were extremely similar to the ROE. My suspicion is that's simply because people look at ROE and

ROA in the same vein. But we didn't see much of a difference between risk measures, whether it's ROA or ROE.

MS. WININGS: So far, we've talked about IRR, ROE and ROA. Each of them has its place, depending upon what type of product and what type of design you have on the reserves that you need to hold. We are now going to move into profit margin. Profit margin is a measure that's often used in life insurance, rather than annuities—where IRRs and ROEs tend to cross over between both product types. The discount rate that is used is typically the pre-tax net-investment-earnings rate. But there can be arguments to use a different rate. I certainly do not intend to suggest one over the other. There are different approaches, and management may dictate what that rate is, as opposed to your pre-tax net-investment-earnings rate.

Essentially, the profit margin is in the discounted profit steams, divided by the discounted premiums. Profits can be either before or after tax, before or after capital, before or after reinsurance. And it will be important that you know which measure you're talking about when you're showing this to someone else. Typically, we see this as a statutory measure. We also see this measure typically for life versus annuities.

One of the reasons that this measure is used is that many companies compare profit margins among each other. You also can use profit margin to compare against your acquisition cost. It doesn't have some of the problems that the IRR has with regard to low reserve issues.

Just looking at it on its own, you don't necessarily reflect when profits occur. As I mentioned earlier, there are a couple of different definitions as to whether or not you include capital, whether or not you look at before or after reinsurance. And then, there isn't a particular metric for every product. It could be some range for term versus whole life, or a spectrum in between.

The example in Table 8 was structured to take a look at a hypothetical product, a 20-year term. I looked at the different ways that you might define the profit margin. I identify how the numerator (or the present value of the profits) is defined, and how the denominator (or the present value of the premium) is defined. In the first column of numbers, the numerator is on a direct basis, and the denominator is gross of reinsurance. This would be the before picture, before you're reinsured. My profit margin is 3.7 percent.

Table 8

Profit Margin Definition - After-Tax			
Numerator	Gross of Reins	Net of Reins	Net of Reins
Denominator	Gross	Gross	Net of Reins
20 Year Term			
PV Profit	80,431	80,482	80,482
PV Prem	2,149,549	2,149,549	969,543
Profit Margin	3.74	3.74	8.30

In the next column we're looking at what reinsurance is providing in terms of my profit streams. In my hypothetical example, I have slightly higher profits.

In the last column, I wanted to consider what kind of profit margin I am earning based on my retained premiums. And of course, the denominator now changes. That causes the profit margins to increase, because your denominator is much smaller. When you are trying to compare what your profit margin might be in some report, it would be important that you know what that measure was, that you were comparing apples to apples.

Table 9 shows the exact same type of analysis, except that it's examining on an after-capital basis.

Table 9

Profit Margin Definition - After-Tax, After-Capital			
Numerator	Gross of Reins	Net of Reins	Net of Reins
Denominator	Gross	Gross	Net of Reins
20 Year Term			
PV Profit	40,817	71,014	71,014
PV Prem	2,149,549	2,149,549	969,543
Profit Margin	1.90	3.30	7.32

How do companies account for capital requirements? In IRR calculations, companies often will change their target surplus by looking at 225 percent or 250 percent of RBC. More commonly, in looking at profit margins, companies tend to look at assumption stress testing. They will look at 10 percent higher mortality, or 10 percent higher expenses, and measure the variability of the profit margin in this manner, rather than looking at RBC. That's partly because many companies look at profit margins on a before-capital basis, so it's not important to them to change the RBC piece to look at the sensitivity.

Companies will also take a look at adjusting their profit measure. Instead of a 2 percent measure, maybe they'll say, "I want this to earn at least a 2.5 percent measure, because I'm a riskier market, or I'm direct response and have some other target issues to deal with." It's less likely that they'll add assumption pads or look at this capital allocation. Again, this is based on a survey that was taken.

There are other things that companies might do to get a sense of the variability of different profit measures. Companies might do a lot of scenario testing—not just a simple set of five or six deterministic views that you consider on every product, but stochastic mortality to look at just a range of things rather than single sensitivities. I certainly understand the resource constraints. If you were to look at every scenario, you'd never get the pricing done by the time you get a product out to market. But there may be some ways to incorporate more analysis than you currently have.

MR. CHUA: RAROC is something fairly new. I don't think that it's widely adopted. There are challenges in applying it to life insurance, simply because of the complexity of applying the technology. But it's a stochastic exercise. You need an assumed distribution of the random variables that you're studying, the mean and the volatility of the random variables and, ideally, a correlation matrix.

One of the keys is recognition that whatever mean and volatility you use is a sample mean and volatility. In arriving at your risk-adjusted income, you make an adjustment for that fact. You run it based on your best estimate, and then you make an adjustment for the fact that those parameters are estimates. You're usually given, by management or by committee, a threshold to achieve. The risk-adjusted capital (very similar to economic capital) has a very high threshold — relative to expected.

The way that the process works is you run it for each random variable, as if they are independent. And then, you take all of that information in aggregate and adjust for the effect of correlation. Here is an example.
Example: Random variable is equity return.

- Normal distribution with sample mean = 10 percent and sample volatility = 16 percent.
- Threshold for sample mean is 90th percentile = 8 percent.
- Threshold for sample volatility is 90th percentile = 20 percent.
- Threshold for economic capital is 99th percentile.

(A) Generate distribution PVFP assuming mean = 10 percent, vol = 16 percent and get average and 99th percentile. [$100, $(900)].

This is by no means extensive, but it gives you a flavor of what you go through to do RAROC. I'm making an assumption that the random variable is equity return. And the normal distribution is being assumed with an expected sample mean of 10

percent and volatility or standard deviation of 16 percent. The company has decided that for purposes of its risk-adjusted income, it wants to get the RAROC within the 90 percent threshold. For economic capital, it wants to hit the 99th percentile.

So with that, you first generate your present value of future profits (PVFP) stochastically assuming your best guess (which is 10 percent and 16 percent) as the parameters. You compute the average and the 99th percentile. I'm showing those to be $100 and -$900 in this example. You then redo the whole exercise, except that you change the parameters slightly.

- Redo (A) twice and get averages, once assuming m = 8 percent and v = 16 percent, then m = 10 percent and v = 20 percent. [$96, $94]
- Risk-adjusted income = 100 + (100-96) + (100-94) = 90.
- Risk-adjusted capital = 100 + 900 = 10000.
- RAROC = 9 percent before covariance adjustment.
- Example is illustrative; other approaches exist.

You would use a mean of eight (which is the 90th percentile) for the sample mean, with an expected volatility of 16 and you get $96. Then you run it again and get the PVFP with a best guess of 10 and in 90th percentile for volatility (which I assume here to be 20 percent), and you get $94. The risk-adjusted income is simply 100 less those adjustments. The risk-adjusted capital is the difference between your 99 percent threshold (which is -$900) and your expected. So you need $1,000, in order to cover all of that. You get a RAROC number of 9 percent—that's before any covariance adjustment. We're talking about a single random variable. But that is just the first step. There's a lot more to it than this, but it gives you a sense of what it takes.

RAROC is very popular in the banking industry. I believe that it actually was started by the banking industry to help them manage their risk and capital. Most bank products tend to be fairly short-term. It's a much easier endeavor for them. When I was with ING we tried doing this. It definitely was much more challenging, particularly when you begin to look into other risk factors—whether that's operational or legal risk that goes beyond the typical factors that you consider in pricing.

It's typically a risk-management measure, and people have not figured out the best way to reflect that in pricing. But given the direction that things are moving (more into financial engineering), I think that this will become more popular. At least, a segment of it will become more popular.

For a lot of companies, particularly European companies, it is an alternative to embedded value, particularly for those organizations that have several different types of business. Some have assets under management and mutual funds, and

some have banking operations. They want to have some sort of a common basis for determining what kind of risk they are taking across the enterprise.

Because of the stochastic nature, the pro is that it lends itself—because of the economic capital calculation or the risk-adjusted capital calculation—to focus on the tail. Typically, those are in the 95th percentile. In some cases, you'll hear the term CTE-95, which essentially is an average of the last 5 percent worst-case scenarios. It allows you to include risk management in pricing, though it continues to be very challenging. And it allows for consistency with other financial institutions. Other than for interest rates and for mortality and equity return distributions, it is probably very subjective. I think that default assumptions have had a lot of work. But with respect to operational risk, legal risk and lapse risk, we typically assume 100 percent correlation. I don't know if that's true or not. The distribution is a very subjective endeavor. It requires a lot of human and computer resources. It's an evolving methodology.

MS. WININGS: One of the emerging measures that has been talked about recently is embedded value, largely driven by the fact that a lot of multinational corporations use that as a primary measure. They are interested in taking a look at the pricing of products with embedded value. It's had widespread use outside of the United States, and because there have been a lot of acquisitions by multinational corporations of U.S. companies, it has come downstream to us.

Embedded value is used for pricing decisions. But it's also used for making decisions with regard to capital allocations, compensation, etc. The measures can vary, depending on if you're using fully allocated expenses or marginal expenses. Right now, we don't see companies using this as a primary pricing measure. We were recently asked to do a quick informal survey of a lot of the top life writers as to different pricing methodology with respect to embedded value and "macro pricing." We found that neither of those methods was in the mainstream, but I do think that is going to be changing. They think that there are some drawbacks. We've already identified drawbacks with some of our more common measures. Embedded value would take care of some of those. Going forward, if you're just pricing the product at issue, you don't always see the potential variances.

There are other measures that companies use to supplement the primary profit measures. Some of these include break-even analysis, taking a look at the surplus strain and sensitivity analysis (and to me that would be more of a deterministic change of different key pricing assumptions). Some companies also use stochastic analysis, which would be more related to changing interest-rate environments or mortality and doing some more widespread sensitivities. You want to understand the key drivers of your product through this analysis. When trying to manage the product, you can understand what would happen if your management says that they want to change something. You'll know if it is a big deal to your return or not. It's good to go through this process during your pricing, rather than just to stop once you get to an answer that seems to be acceptable.

One of the reasons that you would want to go through a stochastic analysis is to better understand your tail risk. I think that in taking a look at the tail risk, we don't—at least on the life side—consider that as much as annuities. And I think that's something that we probably could do better.

Going back to the risk-management survey, they asked companies to identify which pricing measures were the most primary for their product, and to identify the secondary measure. For term life and whole life, we see profit-margin and IRR as the ones listed more often, with profit margin listed more frequently. And then for universal life and variable universal life, it's flipped. I find that it's true in my experience, as well. But other companies are bringing in the secondary measures. ROE is a big factor. Particularly when layering on new business, you want to see level ROEs over time. Break-even analysis is very important. And for some products, such as variable universal life, you might look at more asset-type measures.

MR. CHUA: If I may add a couple of points to the topic of surplus strain, one of the things that you may hear (especially from the marketing organization) is, "Why do you care about the level of the surplus strain if you're getting your appropriate return?" The response that I typically get is that there is an implied concentration risk by putting a large amount of capital into any particular risk. The second is that, especially with annuities, the expected spread may not be achievable or is harder to achieve in this economic condition, particularly if you have a minimum guarantee. So one must emphasize that all we're doing is coming up with an expected number. And a large amount of investment into a single risk may be riskier, though your mechanical calculations show that you are getting the right rate of return.

On measures that are used for annuities for traditional single premium deferred and flexible premium deferred, what we're finding through surveys is that IRR and ROE are the key measures, followed by ROA. And the same thing is true for market-value adjusted (MVA) products and equity-index products. With respect to variable products in particular, ROE is extremely important, and IRR is secondary. The exception to that is with the bonus products. What you see with a variable annuity with a bonus is that there is enough of a product strain that IRR becomes just as important.

With derivative-base benefits, it is typically stochastic analysis. The reason is that a lot of the more popular benefits right now tend to be based on capital-market options. Can the insurance company balance sheet absorb these types of risk? The typical options tend to have to be accounted for on a mark-to-market basis, so there's potential income volatility associated with them. The common approach is to use financial engineering techniques to help solve the problem—if nothing else, to take the whole risk and slice it into pieces. There will be a piece that you believe that you can absorb into the balance sheet, and a piece (typically the tail) that you have to either buy options for or find reinsurance.

MS. WININGS: We're getting to the point in the presentation at which I want to summarize what we've talked about and identify what we think the lessons might be. Not any one of these measures is perfect. You ought to be looking at more than one measure to better understand how your product works and its profitability. You might want to take a look at the IRR and the ROE, so that you understand a product's year-by-year profits, as well as your return on your initial investment. You may want to take a look at your IRR and ROA if you have a product that does not have high reserves. Or you might want to take a look at your profit margin and your stat IRR. And by and large, the clients that I work with do look at more than one measure. It's very important to consider.

It's important that you understand what your measure is and what it is not—exactly what's included in it so that when you're having a conversation with somebody, you know that you're on the same wavelength, and you understand what you're trying to measure. Understand what your key drivers are, so that as things change going forward, you have a sense of how that might change your returns.

As we listen to the evening news every night, we hear about earnings from different companies and how that drives the stock market. Earnings are a big consideration for people who are investing. Consider whether or not investors really are looking at these earnings versus a present-value measure. What direction might we be going in the future in terms of measuring profitability? If you ponder the examples that we gave today, they're mostly at-present-value calculations, as opposed to earnings that are coming out year by year. But that seems to be a big driver, according to all of the analysts. Is that going to change how our management looks at how we develop products? How will we handle that, and will our pricing measures change?

FROM THE FLOOR: As a multiline company that has term insurance, traditional life, variable life, annuities, banking products and only a finite amount of surplus, how should we decide how to invest our surplus—in an annuity product, or term-insurance product, etc.? Are we going to encourage sales of our banking products versus our mutual funds? How do we deal with that, given that our profit measures are different? How do you make the decision when you're sitting with senior management and they're trying to decide whether to work on an annuity product or a traditional life product? How do we help them understand which one to choose?

MR. CHUA: My opinion is that the common thread is figuring out the right amount of investment or capital that you should put into the product, as opposed to what the rating agency or the regulatory environment would indicate. I'm aware of at least one organization that has done such a good job with their asset/liability modeling (ALM), that one of the components in their capital requirements is lower than prescribed. So I think that in order to come to a common ground, the common thread is figuring out the appropriate investment consistent with the risk, and then perform an IRR calculation.

MS. WININGS: I think that companies that deal with that often have to come up with some sort of comparison point. And the method that a lot of companies would use for that would be a shareholder value. That was the terminology in my prior stock company experience. But I think that embedded value would be another terminology for that. If you have some common metric for each one of those lines of business, you could then do a comparison to see which one would provide more value than the other. That's the way that some companies handle that decision. I don't know that all companies take the time to step back and look at that when they are pricing a particular product. They might say that they are going to allocate a certain amount once a year. Then when it is time to price that product, they'll try to find that embedded value or a return with that. But I think that all companies grapple with trying to come up with some measure to make that relative decision.

FROM THE FLOOR: Early on, when you were speaking of the cons for the IRR, you said that it may be inappropriate if investment is inconsistent with risk. Could you elaborate a little more?

MR. CHUA: If you put money into in an investment-grade bond, the market tells you the appropriate rate of return that you're entitled to. If you go down the quality ladder, you get an expected return that has a higher risk premium. My comment refers to the fact that we use either the rating agency or the regulatory capital as the basis for pricing the product. And it's not clear that that amount of the investment or capital is always an appropriate level for the product. If that is not the case, then the same IRR that you're getting for product A versus product B may not be appropriate. You're maybe getting the same rate of return holding the same capital, but one product may be riskier than the other.

One example might be to examine a high-strained product, with a significant first-year strain. Expected IRR may be the same as for a product with a lower strain. In my opinion, the higher-strain product, because it demands a higher spread, actually has a higher risk.

FROM THE FLOOR: My life insurance company has a mutual fund. I'm curious about your experience with applying actuarial techniques on long-range profit measures to mutual funds. Do you see companies doing that? We have the same sort of situation with the first question, making decisions about use of capital on mutual funds versus insurance products.

MR. CHUA: Take C-shares with minimal acquisition cost and significant fixed-infrastructure cost as an example. One approach for arriving at an appropriate ROA level is to measure the income volatility associated with the mutual funds and set the spread requirements so that the fixed infrastructure costs are covered 90 percent to 95 percent of the time. Comparing the income volatility in fixed annuities against variable annuities and mutual funds may provide some insight. Another approach is to determine the cost of hedging away the income volatility risk. The hedge cost together with an appropriate return and recoverability of the

infrastructure cost can then be converted into an appropriate spread on assets. The challenge is finding an appropriate hedging instrument with a high correlation to the mutual fund under study. A lack of hedge effectiveness may be a waste of money and resources.

Chart 1

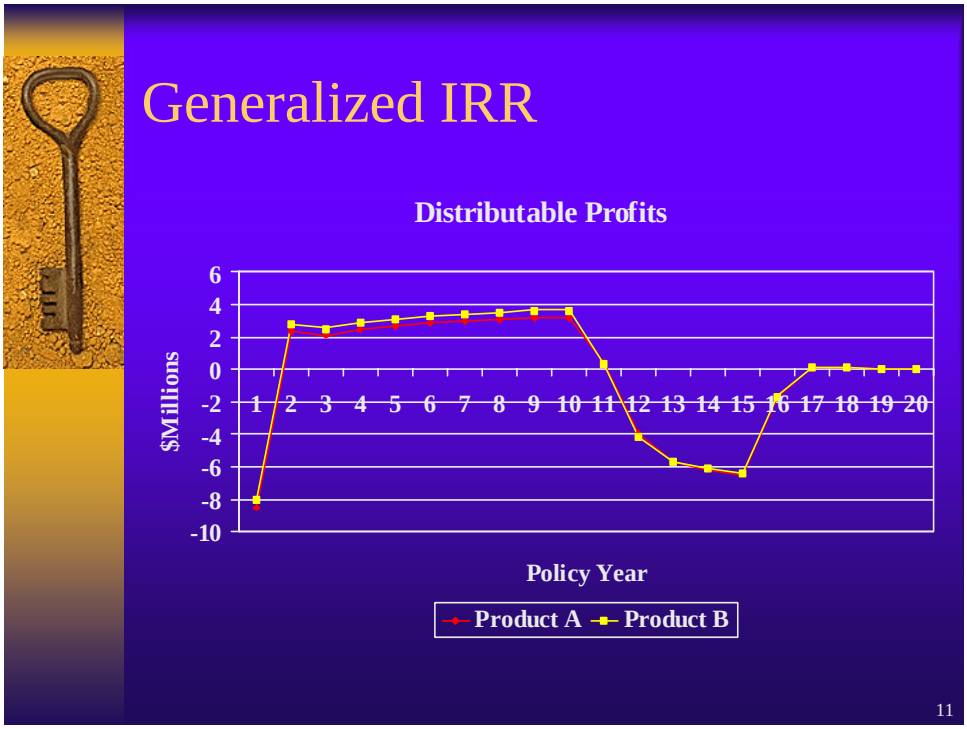
Generalized IRR
Distributable Profits
$Millions
6
4
2
0
-2
-4
-6
-8
-10
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
Policy Year
Product A
Product B
11