RECORD OF SOCIETY OF ACTUARIES
1995 VOL. 21 NO. 4B

DEVELOPING OPTIMIZED INVESTMENT STRATEGIES FOR LIFE INSURANCE COMPANIES

Instructors:	STEPHEN D. REDDY
	GREGORY J. ROEMELT
Co-Instructor:	MICHAEL WALTER PESKIN
Recorder:	GREGORY J. ROEMELT

Optimization is frequently used with pension liabilities. How can these concepts and tools be applied to life and annuity liabilities?

MR. GREGORY J. ROEMELT: I'm a consultant at Tillinghast in Chicago. We're privileged to have Steve Reddy from Morgan Stanley, who will discuss using derivatives to help optimize the insurance company portfolios, and Michael Peskin of Michael Peskin & Associates, who will discuss using some of the pension theory in developing optimized portfolios in an insurance company environment.

MR. STEPHEN D. REDDY: The topic is optimizing—developing optimized strategies for life insurance companies. When Greg asked me about possibly speaking here, I had an interest because we had done some work along these lines that involved derivatives. So I had to check with him to make sure that was acceptable. I think that derivatives have gotten a lot of negative press in the last couple years or so, primarily because they've been used in the wrong way in many cases, and those are the cases that got into the headlines. But there is a proper way to use derivatives, and I will show you examples of the right way to use them or *a* right way to use them, but certainly not the only way.

We're going to try to do some optimizing of profits, and I guess the question is, what does optimizing mean? Generally it means maximizing something such as earnings or profits subject to some constraints, and what those constraints are remain to be defined. But, generally speaking, that's what profit optimization is. I guess there are two ways of going about optimizing earnings in a life insurance company. One would be to attempt to develop an investment strategy that did maximize earnings subject to some specific constraints. I think that's actually more difficult than taking an existing portfolio and projecting earnings out, seeing where the peaks and valleys are, and then, through an overlay strategy with derivatives, smooth out the peaks and valleys. A hedging program should take away from some of the good results and apply them to the situations when earnings are not as good and cause problems.

Anyway, I will describe a six-step process to go about optimizing earnings with derivatives, and this is something that you ought to be able to take back and apply in your shop. Now, maybe all the details won't be perfectly clear, but the process ought to be clear. To the extent details aren't clear, I'd be happy to talk more about that after the session. But this is something that you should be able to take back, and I hope you'll agree that the approach makes a lot of sense.

Essentially we're talking about some systematic way of using some derivatives in an optimization process so the end result is a hedging program in which we actually are impacting the scenarios we want in the way that we want. We take away earnings from those scenarios where we don't really need them and shift them over to those cases where we do need them. We're going to do this by using some simple derivatives, interest rate swaps, caps, and floors. They are your basic building blocks (and a relatively small set) of derivatives. By using those we will be shifting some profits from certain scenarios to other scenarios and maybe from certain years to other years, and we can do this across any projected set of interest rate paths.

What is the process that I'm talking about here? Basically, the six steps are as follows. One is you determine some set of interest rate paths upon which you wish to evaluate your company. Now, this could be as simple as taking the New York 7, which probably isn't the ideal case, but that's a situation that most companies will have already through cash-flow testing. So you could actually take those seven and do some optimization based on that, or even apply some weighted averages to those seven, but ideally you do some other set with more than seven paths, up to 20, 40, 50, 100, or whatever it might be. Ideally, they're equally weighted, but they wouldn't have to be. You could apply weights to whatever set of paths you come up with. So whatever interest rate generator you happen to be employing, it doesn't matter. Whatever set of interest rate paths you want to measure your company's earnings or the volatility of earnings on, can be used in this process.

The second thing is to project profits or earnings along each of those interest rate paths, and it can really be for any number of years. I've specified five years, because I think when going beyond five years you start to lose confidence in the actual earnings beyond. Management generally will manage to earnings within a five-year period.

Once you have those two things, the next step would be to specify some universe of hedging instruments that you'd be willing to consider to buy or sell to accomplish the hedging and the redistribution of earnings along these various interest rate paths.

Once we have that, the next two steps are to specify the constraints that we're going to impose on the situation, and an example of a constraint would be we never want earnings to drop below $5 million in any calendar year over the next five years. Obviously, setting constraints makes sense. To set the constraints you'd need to look at where earnings are projected in the unhedged situation. The distribution of earnings may range from -$10 million to $70 million. That may tell you that a reasonable level to shoot for is no less than $10 million in any particular year but, again, that's judgmental. You'd need to look at the unhedged projections to determine where an acceptable set of constraints might lie. You can define constraints for each and every calendar year, or the constraints can be defined on a cumulative basis.

Once you define some constraints, then you need to determine one objective function that you wish to maximize or minimize, depending on what it is, subject to meeting the constraints you specified up above. Now, it's possible you might define a set of constraints that is so burdensome that there is no solution, and you can't buy or sell any combination of hedges to satisfy the constraints. If you can't satisfy the constraints, satisfying an objective function is meaningless. Constraints have to be satisfied first. So, in the example I mentioned before, the unhedged earnings range from -$10 million to $70 million. You want to apply some hedges so that the earnings are never less than $70 million. In any case you won't be able to do that. That's fairly obvious. Anyway, subject to constraints, you could specify some objective function, which might mean to maximize the cumulative earnings through a five-year period, or the mean earnings over all the interest rate paths that you're running.

Finally, once you've gone through the first five steps, the sixth step is to essentially find the answer. The way we're doing that is through linear programming techniques. Through linear programming you can find an optimal solution, which means maximizing the objective functions subject to the specified constraints, buying and selling only the specified instruments that are specified in Step 3. You get some very interesting solutions, depending upon what the constraints and objective functions are, as well as obviously the input specified in Steps 1 through 3.

Regarding examples of constraints and objective functions that could be used in this process, one is the notional amount of any particular hedging instrument. You might say, well, depending on the size of your company and perhaps management's lack of comfort with derivatives or whatever, you may decide that you don't want to buy any more than $100 million of any particular interest rate cap or swap. That could be specified as a constraint.

On the other hand, you may decide that you don't want any more than $500 million of all hedging instruments combined. That constraint could be specified on either an absolute basis or on a net basis. In some cases you might be buying interest rate caps and selling floors to generate income. In that case you might consider just the net notional amount, but those can be specified either way. You might also want to constrain the amount of up-front premium that you're paying for any of these instruments. Interest rate swaps don't have any up-front premium. Therefore, there's no immediate effect on the balance sheet or income statement, whereas, when buying an interest rate cap there would be. You'd be booking an asset and spending money to acquire that asset.

You may want to specify some limit on the negative carry for an interest rate swap, which is the difference between the fixed rate and the initial floating rate. In terms of earnings you could constrain the minimum earnings, meaning the minimum over all the interest rate paths that you've projected, such that the earnings never get less than x in any of those paths. Or you might want to constrain the mean of the earnings so that the mean is never less than x over any of those paths. Or you might want to constrain the range of earnings, not allowing the difference between the best path and the worst path to get greater than $40 million, for example. If the current range is $80 million, maybe you want to constrain that to be no greater than $40 million.

The final example, the mean absolute deviation, is sort of a proxy for standard deviation. We can't use standard deviation here because that involves a quadratic, and the linear program won't be able to solve that, but this is a similar concept. Again, there are many other examples to things you could constrain, and anything you can constrain can also be used in your objective function.

Again, as we get into some examples here, I hope this will become clearer and make a little more sense. So, let's get into the first example. Table 1 shows step 1, which is selection of interest rate paths. I actually have four examples or case studies in here. Each uses the same set of interest rate paths, projected earnings, and set of hedging instruments, and there are 20 interest rate paths in this example. I have only shown a few of the paths, but I think you get the idea. Twenty paths go across. Interest rates are projected out five years on a quarterly basis, and only the five-year Treasury rate for each of those paths is shown. So you can see it starts out as 6.94 in each case and then branches out, depending upon the path. So those interest rate paths will be used to evaluate the set of earnings that we have for this hypothetical company.

TABLE 1
STEP 1: SELECTION OF INTEREST RATE PATHS
FIVE-YEAR CONSTANT MATURITY TREASURY RATES (PERCENTAGE)

Quarter	Path 1	Path 2	Path 3	. . .	Path 18	Path 19	Path 20
1	6.94	6.94	6.94	. . .	6.94	6.94	6.94
2	7.19	8.80	6.44	. . .	6.04	6.90	7.37
3	6.24	9.36	6.61	. . .	6.00	6.52	7.15
.	.	.	.	. . .	.	.	.
.	.	.	.	. . .	.	.	.
.	.	.	.	. . .	.	.	.
18	5.68	9.33	7.76	. . .	5.09	7.98	7.48
19	5.96	9.20	8.44	. . .	5.21	8.62	7.11
20	5.67	9.86	9.17	. . .	5.15	8.90	7.26

Step 2 looks similar. Table 2 is the projection of earnings along each of those interest rate paths that we just looked at. So, again, there are 20 paths, 20 sets of earnings along 20 quarters, as it turns out, five years of projected earnings, and these earnings happened to come from a hypothetical single premium deferred annuity (SPDA) block. So they're not totally made up. They actually represent projected earnings off some interest-sensitive block of liabilities.

TABLE 2
STEP 2: PROJECTION OF EARNINGS ALONG EACH OF THE INTEREST RATE PATHS QUARTERLY EARNINGS (IN $1000S)

Quarter	Path 1	Path 2	Path 3	. . .	Path 18	Path 19	Path 20
1	$4,687	$4,690	$4,970	. . .	$4,830	$4,865	$4,792
2	4,915	3,843	4,952	. . .	5,340	5,107	4,715
3	5,044	3,344	5,118	. . .	5,593	4,985	4,404
.	.	.	.	. . .	.	.	.
.	.	.	.	. . .	.	.	.
.	.	.	.	. . .	.	.	.
18	4,696	2,214	3,241	. . .	5,404	2,390	2,910
19	4,122	1,956	2,673	. . .	4,689	1,938	2,381
20	$3,792	$2,003	$2,433	. . .	$4,323	$1,745	$2,436

See footnote after Table 1.

Obviously, you can't see all the quarters and the paths, but of the ones that are shown, Path 2 shows the lowest earnings, I believe, or at least in certain years Path 2 and Path 19, I guess, are the worst of the two. Path 2 did have the highest interest rate spike, particularly early on, and Path 19, I guess, was also one of the worst paths. So you can see some evidence here of the interest sensitivity of the block and earnings falling off in the higher interest rate scenarios.

We've projected earnings along those paths. Now in Step 3 (Table 3) we define some universe of hedging instruments. Again, what I have here is just a small set of interest rate swaps, caps, and floors. The swaps involve just paying a fixed rate and receiving the five-year constant maturity treasury (CMT) rate over some period. The period is in the next-to-the-last column on the right. That's the term of the swap. So there is a three- or five-year swap there. What you see in the second and third columns is essentially the bid as spread for each of these instruments. In this case, if the insurance company wants to receive a floating rate, it would pay 0.0758 and receive the five-year CMT over the next three years on a quarterly basis. If it wanted to receive a fixed rate and pay floating, or receive 0.0748 and pay the five-year CMT.

Similarly, all the way down the line for the caps and floors you see the up-front premium for buying, for example, 0.0497 for a five-year cap, which is an at-the-money cap. The strike price is 0.0694. Going down in 50-basis-point increments to a much less expensive cap, 87 basis points up front for a five-year cap, that's 250 basis points out of the money. The strike price is 0.0944.

In the actual examples I use both three- and five-year caps and floors, but I just showed here the five-year caps and floors. But there are also three-year caps and floors at the same strike prices as the five-year caps and floors, and again just the two interest rate swaps. We're going to consider a fairly small set of instruments in the hedging process here. These are the only things we're going to buy or sell. Unless we constrain it otherwise, you would be allowed to buy or sell any or all of these instruments.

TABLE 3
STEP 3: SELECTION OF POTENTIAL HEDGING INSTRUMENTS

Instrument	Premium-Buy	Premium-Sell	For (Years)	Strike
Swap	0.0758	0.0748	3.00	0.0000
Swap	0.0775	0.0765	5.00	0.0000
Cap	0.0497	0.0436	5.00	0.0694
Cap	0.0368	0.0311	5.00	0.0744
Cap	0.0264	0.0214	5.00	0.0794
Cap	0.0188	0.0144	5.00	0.0844
Cap	0.0127	0.0093	5.00	0.0894
Cap	0.0087	0.0060	5.00	0.0944
Floor	0.0163	0.0126	5.00	0.0694
Floor	0.0097	0.0067	5.00	0.0644
Floor	0.0049	0.0030	5.00	0.0594
Floor	0.0018	0.0009	5.00	0.0544

See footnote after Table 1.

On to Steps 4 and 5, which are defining the constraints. In this case, and again this will become a little more meaningful when we go forward, we decided to not allow any particular instrument to have a notional amount greater than the $500 million. The total absolute value of all notional amounts is no greater than $1 billion. So there is an individual limit and a total limit for the hedging instruments, and then we're constraining earnings in each of the five years.

In this case it shows to set the constraints on a calendar-year basis even though the projected earnings were quarterly. It could have been done. You could have set the constraints on a quarterly basis as well, but management's more likely to want to manage on a calendar-year basis; that's what we've done here.

These constraints were set based upon when you look at the range of unhedged earnings. These constraints were selected to bring up the worst case to some level that was deemed acceptable. Seven specific constraints are specified. The objective would be to, assuming we can meet those constraints, maximize the cumulative present value of earnings through five years, discounted at 10%. So let's see what it takes to do that.

CASE STUDY 1

Step 4: Selection of Constraints*	
Constraint 1:	Notional amount for each hedging instrument ≤ $500,000,000
Constraint 2:	Total absolute value of all notional amounts ≤ $1 billion
Constraint 3:	Year 1 earnings ≥ $12,500,000
Constraint 4:	Year 2 earnings ≥ $10,000,000
Constraint 5:	Year 3 earnings ≥ $10,000,000
Constraint 6:	Year 4 earnings ≥ $7,500,000
Constraint 7:	Year 5 earnings ≥ $7,500,000
Step 5: Selection of Objective Function	
	Maximize the cumulative present value of earnings through year 5, discounted at 10%

*See footnote after Table 1.

Table 4 is actually the answer that the linear program produces. This is the actual answer in that it tells you the actual amounts or specific amounts of each instrument you would need to buy or sell to achieve those constraints or satisfy the constraints and maximize the objective function. You can see, for example, that the first two constraints were met and none of the notional amounts exceeds $500 million, either plus or minus $500 million, and the total absolute value of all the hedging instruments actually did exactly equal $1 billion.

That constraint got hit and actually came into play here. The absolute value of all those hedging instruments, the five things shown up there, does add up to $1 billion. So that constraint came into play. Basically, the program said, we need to buy an interest rate swap here, $134 million, buy a cap at the money, $370 million, but then sell a cap, a slightly larger amount, 250 basis points out of the money.

Of course, selling that cap generates some income which is needed to help finance this whole thing, as does selling a floor. Selling the floor was done because we could afford to generate income there. We did need protection in the cases where the insurance company would have to pay off the floor, which obviously would be a lower interest rate scenario.

TABLE 4
STEP 6: THE SOLUTION

Instrument	Notional	Unit Premium	Upfront ($)	For (years)	Strike
Swap	$134,338,000	0.0758	0	3.00	0.0000
Swap	1,677,000	0.0775	0	5.00	0.0000
Cap	370,971,000	0.0497	$18,437,000	5.00	0.0694
Cap	(408,486,000)	0.0041	(1,675,000)	3.00	0.0894
Floor	(84,528,000)	0.0067	$ (566,000)	5.00	0.0644
Net Total:	$13,972,000		$16,196,000		
Absolute Total:	**$1,000,000,000**				

See footnote after Table 1.

Let's see what we actually accomplished by doing all this. On the left-hand side of Table 5, you see the earnings profile in the unhedged situation. Looking at the unhedged situation, the worst case was $14,323,000, year one, and then it went negative in years four and five. You can also see the range of earnings, the best case in each of those cases, as well as the means. Well, that's where we pulled out the constraints to specify here.

TABLE 5
SUMMARY OF EARNINGS* (IN $1000s)

	Unhedged			Hedged		
	Mean	**Low**	**High**	**Mean**	**Low**	**High**
Profit:						
Year 1	$18,488	$14,323	$22,259	$17,510	$14,328	$20,966
Year 2	15,149	3,058	28,109	14,907	10,000	21,382
Year 3	13,579	1,489	24,011	13,287	10,000	16,812
Year 4	11,864	(456)	21,252	13,322	8,764	16,821
Year 5	10,386	(455)	19,504	12,579	7,500	16,980
Cumulative Present Value:						
Year 1	$17,647	$13,744	$21,189	$16,704	$13,683	$20,012
Year 2	30,787	16,409	44,117	29,639	25,438	33,624
Year 3	41,489	17,585	63,047	40,109	35,107	46,471
Year 4	50,003	17,255	77,848	49,661	43,418	55,871
Year 5	56,778	16,959	89,814	57,860	50,384	65,125

*Straight line amortization of caps and floors has been assumed here. Other accounting methods could also be used.
See footnote after Table 1.

Again, we specified a worst case, I think $12 million, $12.5 million the first year, $10 million in the second and third years, $7.5 in the fourth and fifth years. Now moving over to the right-hand side you can see that the program actually hit the constraints in three of those five years, which is actually evidence that the program worked. It did what you told it to do, which is kind of reassuring.

If the linear program spits out an answer, that's one thing, but to actually take the results and then add the results of those hedging instruments to the unhedged profits gives you reassurance when you see that the combined results actually hit some of the constraints. Basically, it tells you that it did what it was supposed to do. So, we actually were successful in raising our worst case results in four of the five years.

We satisfied the constraints everywhere, and the thing that it actually maximized, subject to those constraints, is the other highlighted number down there, the cumulative present value (PV) of earnings through five years. According to the linear program, this combination of instruments maximizes that figure subject to all constraints that were specified.

Let's move on to the second case where we're going to have just a slight variation from the first example. Here, the seven constraints we specified are identical. So we're not going to change those at all. But subject to those constraints, instead of maximizing the present value of earnings, maybe management is reluctant to pay too much, at least up front, to implement this hedging program. Management might specify, well, let's see what we can do paying as little up front as possible. We could, as an objective function, minimize the total up-front premium required to be paid to satisfy those constraints.

CASE STUDY 2

Step 1–3:	Same as Case Study 1	
Step 4:	Selection of Constraints	
	Constraint 1:	Notional amount for each hedging instrument ≤ $500,000,000
	Constraint 2:	Total absolute value of all notional amounts ≤ $1 billion
	Constraint 3:	Year 1 earnings ≥ $12,500,000
	Constraint 4:	Year 2 earnings ≥ $10,000,000
	Constraint 5:	Year 3 earnings ≥ $10,000,000
	Constraint 6:	Year 4 earnings ≥ $7,500,000
	Constraint 7:	Year 5 earnings ≥ $7,500,000
Step 5:	Selection of Objective Function	
	Minimize the net total up-front premium required of the counterparty for all caps and floors in the hedging program.	

See footnote after Table 1.

Table 6 shows a different answer, which allows some similarities but some clear differences as well. It still meets the total absolute value constraint of $1 billion, but the up-front premium, which was what we were trying to minimize came down to $2,691,000, whereas it was $16 million in the first case. Clearly it appears evident that the linear program did what it was supposed to do or certainly it gives the appearance that it did.

Moving forward to Table 7, the numbers show both the unhedged and hedged results. The unhedged, again, is the same set of numbers you saw before. On the hedged case, we're hitting the constraints in three of the five years, but it's actually a different three years, the first, third and fifth instead of the second, third, and fifth. The mean earnings figure is now lower than it was in the first example because we're not optimizing this one anymore. We chose to optimize reducing the amount of premium you needed to pay, but it's still not a bad result. You can't really tell just by this in isolation, but if you were to compare this graph with the one we looked at in Case 1, and actually held them up to each other, there's actually very little difference. These two cases are fairly close. The constraints were identical in each case, and those became a large part of the solution. But the linear program specified a different set of results.

TABLE 6
STEP 6: THE SOLUTION

Instrument	Notional	Unit Premium	Upfront ($)	For (years)	Strike
Swap	$218,121,000	0.0775	$0	5.00	0.0000
Cap	116,580,000	0.0497	5,794,000	5.00	0.0694
Floor	124,087,000	0.0009	112,000	3.00	0.0544
Cap	(158,216,000)	0.0041	(649,000)	3.00	0.0894
Floor	$(382,996,000)	0.0067	$(2,566,000)	5.00	0.0644
Net Total:	$(82,424,000)		$ 2,691,000		
Absolute Total:	$1,000,000,000				

See footnote after Table 1.

TABLE 7
SUMMARY OF EARNINGS* (IN $1,000S)

	Unhedged			Hedged		
	Mean	Low	High	Mean	Low	High
Profit:						
Year 1	$18,488	$14,323	$22,259	$17,495	**$12,500**	$20,224
Year 2	15,149	3,058	28,109	14,646	10,218	21,945
Year 3	13,579	1,489	24,011	13,016	**10,000**	16,234
Year 4	11,864	(456)	21,252	12,375	7,801	16,645
Year 5	10,386	(455)	19,504	11,797	**7,500**	16,637
Cumulative Present Value:						
Year 1	$17,647	$13,744	$21,189	$16,699	$11,997	$19,297
Year 2	30,787	16,409	44,117	29,407	25,623	32,752
Year 3	41,489	17,585	63,047	39,660	35,551	45,131
Year 4	50,003	17,255	77,848	48,536	42,859	54,281
Year 5	56,778	16,959	89,814	**56,223**	49,456	62,715

*Straight line amortization of caps and floors has been assumed here. Other accounting methods could also be used.
See footnote after Table 1.

Looking back on it, more interest rate swaps were used because they didn't involve an up-front premium. A greater amount of caps and floors were sold to generate income, and essentially that generated incomes and lowered the amount of premium that had to be put up front. It also reduced slightly the amount of earnings that the company would achieve in a very high interest rate scenario where one of the caps was paying off. Well, in this second solution here we sold more of the cap. Therefore, we'd actually start paying off more in that particular very high interest rate scenario. But given the set of paths that were projected, that was an acceptable result.

We also sold a lot more, four in the first case, out of the money. We actually sold a floor, a large amount of floor, $382 million versus $84 million in the first case, but then we bought a floor 100 basis points of the money to cap our exposure there, but again the combination of those two things generated a little income, which helped to achieve the objective that we had specified. Again, this will look very similar to the first case; there are very slight differences.

Were going to change things slightly. In the third case we have only one constraint. The specific limit for any hedging instrument is $500 million. There are no other constraints in terms of earnings by year or the total absolute value of the notional amounts.

Our objective function will be to maximize the lowest annual earnings for any year or any path. This is basically saying, I want to look at the worst case—the worst projected earnings in any year in any path over the next five years—and I want to bring that up to the highest level possible.

CASE STUDY 3

Step 1–3:	Same as Case Study 1
Step 4:	Selection of Constraints
	Constraint 1: Notional amount for each hedging instrument ≤ $500,000,000
Step 5:	Selection of Objective Function
	Maximize the lowest annual earnings for any year and path.

See footnote after Table 1.

Here again is the solution that was produced (Table 8). What we have here is somewhat different. We now exceed the $1 billion in total notional amount, up to $1,654,320,000 instead of $1 billion. We're also spending $19 million up front. That's not a $19 million hit to earnings. A check is written, and an asset is booked at that time.

TABLE 8
STEP 6: THE SOLUTION

Instrument	Notional	Unit Premium	Upfront ($)	For (years)	Strike
Cap	302,599,000	0.0265	$8,019,000	3.00	0.0694
Cap	199,592,000	0.0497	9,920,000	5.00	0.0694
Cap	471,913,000	0.0127	5,993,000	5.00	0.0894
Floor	(248,891,000)	0.0126	(3,136,000)	5.00	0.0694
Floor	(431,325,000)	0.0030	$(1,294,000)	5.00	0.0594
Net Total:	$293,888,000		$19,502,000		
Absolute Total:	$1,654,320,000				

See footnote after Table 1.

It is an accounting entry or something to be cognizant of, and it's something that management may or may not want to avoid. Anyway, the solution here involves buying more caps than in the previous cases, and also selling more floors. The profile is quite different in terms of the solution.

Let's see what was achieved when we did that. Again, the left-hand side of Table 9 is identical to what we saw before. That's the unhedged result. The right-hand side, the $8,610 as it turned out, was the highest that the earnings could be raised to in the worst case out of all five years, and there were 20 paths in this case. So 100 path years were being optimized in this particular instance, and $8,610 was as high as the worst case could be brought up.

That actually is the optimal solution in this case, given the specifications. That's relative to the previous examples you had seen. The low earnings in year two, for example, were $10 million. So we've relaxed that constraint in this particular example for the purpose of raising up the lowest earnings in the subsequent years, the fourth and fifth years in particular.

TABLE 9
SUMMARY OF EARNINGS* (IN $1000S)

	Unhedged			Hedged		
	Mean	**Low**	**High**	**Mean**	**Low**	**High**
Profit:						
Year 1	$18,488	$14,323	$22,259	$16,729	$10,279	$30,744
Year 2	15,149	3,058	28,109	15,127	**8,610**	35,919
Year 3	13,579	1,489	24,011	13,823	8,734	28,919
Year 4	11,864	(456)	21,252	12,748	**8,610**	20,196
Year 5	10,386	(455)	19,504	12,293	**8,610**	20,795
Cumulative Present Value:						
Year 1	$17,647	$13,744	$21,189	$15,952	$9,884	$29,103
Year 2	30,787	16,409	44,117	29,078	21,058	60,294
Year 3	41,489	17,585	63,047	39,955	30,359	83,106
Year 4	50,003	17,255	77,848	49,102	38,085	94,956
Year 5	56,778	16,959	89,814	57,113	46,520	100,566

*Straight line amortization of caps and floors has been assumed here. Other accounting methods could also be used.
See footnote after Table 1.

Again, the constraints in case 4 are identical. There's only one constraint. That's the amount of any particular hedging instrument. The objective here, rather than trying to optimize or raise up the worst result in any path or year, is to raise up the mean of earnings for any year.

That's the mean earnings across all paths. Obviously, that allows for some distribution in any year, but we wanted to maximize the mean earnings for any year across all paths in those years.

CASE STUDY 4

Step 1–3:	Same as Case Study 1
Step 4:	Selection of Constraints Constraint 1: Notional amount for each hedging instrument ≤ $500,000,000
Step 5:	Selection of Objective Function Maximize the lowest annual mean earnings for any year.

See footnote in Table 1.

Let's see what it did there. Table 10 shows we actually are buying and selling more instruments here. We're up to $1,985,952,000 of notional amount and actually $62,227,000 in terms of up-front premium, which is much higher than before. We are buying four caps, which is more than before, and let's see what that does.

TABLE 10
STEP 6: THE SOLUTION

Instrument	Notional	Unit Premium	Upfront ($)	For (years)	Strike
Cap	$500,000,000	0.0265	$13,250	3.00	0.0694
Cap	426,193,000	0.0186	7,927	3.00	0.0744
Cap	500,000,000	0.0497	24,850	5.00	0.0694
Cap	440,217,000	0.0368	16,200	5.00	0.0744
Swap	$(119,542,000)	0.0765	$0	5.00	0.0000
Net Total:	$1,746,868,000		$62,227		
Absolute Total:	$1,985,952,000				

See footnote after Table 1.

Again, instead of modifying or improving the results in the low earnings paths, we've brought up the mean column, which is what we intended to do in this case (Table 11). We've brought it up to $15,065 in three of the five years, and looking over at the unhedged case, in the left-hand column, you can see that those numbers were, in the third, fourth and fifth years, $13,579,000, $11,864,000, and $10,386,000.

We were able to improve those through this particular objective function. In the process we let the worst case result slip somewhat. They dropped down to the $1–2 million range in the second and third year, which is obviously much worse than the previous examples.

An income-oriented approach is something that's very easy to grasp just intuitively. If you don't use that and just use sort of an economic surplus kind of approach, it's a little harder to tell what the impact is on an income statement perspective or from an income statement perspective. You may have an idea that if you buy a certain cap, you may improve your economic surplus by a certain amount if interest rates spike up 200 basis points. But you don't know how that will unfold over time and how it will affect the income statement during the next three years, for example. This allows you to see the effect on the income statement in each period going forward.

TABLE 11
SUMMARY OF EARNINGS* (IN $1000S)

	Unhedged			Hedged		
	Mean	**Low**	**High**	**Mean**	**Low**	**High**
Profit:						
Year 1	$18,488	$14,323	$22,259	$15,065	$6,274	$48,719
Year 2	15,149	3,058	28,109	16,201	1,972	69,072
Year 3	13,579	1,489	24,011	15,065	1,707	57,765
Year 4	11,864	(456)	21,252	15,065	5,700	28,537
Year 5	10,386	(455)	19,504	15,119	4,209	28,938
Cumulative Present Value:						
Year 1	$17,647	$13,744	$21,189	$14,309	$5,975	$45,934
Year 2	30,787	16,409	44,117	28,370	11,092	105,898
Year 3	41,489	17,585	63,047	40,220	13,545	151,459
Year 4	50,003	17,255	77,848	51,012	19,619	169,440
Year 5	56,778	16,959	89,814	60,863	24,657	179,475

*Straight line amortization of caps and floors has been assumed here. Other accounting methods could also be used.
See footnote after Table 1.

There's more reliance on marking to market or unwinding the positions, which is one reason why companies are afraid to buy options or caps. To the extent they have to go back to the investment bank to have it buy them back, they may think that they won't get good pricing in trying to unwind the position. This doesn't involve any of that. It's just buying and holding the instruments, basically.

Knowing the price up front allows you to do the analysis, and it's not dependent upon unwinding the positions or marking them to market. Time decay is explicitly recognized. What I mean by that is that the passage of time is being explicitly accounted for in the analysis here. We're looking at the earnings as they unfold over time.

Finally, the process is easy to modify or update over time. As you go out one year into the future, you could again repeat the analysis and perhaps just layer on additional instruments to deal with the out years as another year or two passes.

If you've taken and gone through this process with one set of interest rate paths, you can actually take the solution set and then measure the results—how well the solution set works against some other set of interest rate paths. So if you're not totally confident in one set of paths, the solution that is derived here could be tested against other sets of paths.

Alternatively, you could have the linear program solve for the set of solutions based upon multiple sets of paths. In a way it's like optimizing against more paths to start with rather than just the 20 that were used here. You could hedge against multiple profit projections.

Let's say you weren't totally confident in one set of projections. Maybe you want to test an alternative crediting strategy or alternative surrender behavior that might result. Again, you could optimize against those two alternative sets of paths or optimize against one and test it against another or vice versa. You've seen that this process is a way of judging the merits of buying and selling certain amounts of caps. So it's a systematic way of evaluating swaps and certain options.

If you think volatility's will be higher, project a set of interest rate paths based on higher volatility. Test the results against that or find a solution set against that higher volatility set or vice versa. This process could be done over different horizons, and the results compared for each of those different solution sets and judgments could be made about what the best horizon is to work against.

MR. MICHAEL WALTER PESKIN: While I was at Morgan Stanley, which I left two-and-a-half years ago to start my own firm, I developed a system for pension plans, that had a kind of top-down corporate finance focus. I looked at the pension plan as being a collateralized bond obligation of the corporation.

The corporation's objective was to minimize the present value of future contributions minus the call on surplus that it owned. Corporations were able to save enormous amounts of money, but they needed to change, to a large extent, what they were doing, particularly with asset allocation. The more I thought about it, the more convinced I became that the same kind of top-down corporate finance approach would also have applicability in insuranceland. There were three reasons for this.

I was aware personally that some insurance companies don't pay any, or very little attention to asset/liability management, that they're forced to by regulations. But those that did pay attention tended to pay attention on a product line basis, putting a lot of attention to getting some product line in detail, accurate in cash-flow matching, but the aggregate asset allocation was really a sum of these details. Some you couldn't match at all, and the aggregate was there for probably overconstraint. I was aware that insurance liabilities had a lot of noise. What I mean by that is even if you go to your lowest-risk-possible portfolio and you do the best job you can of matching liabilities, you cannot match them.

Many liabilities have little to do with the capital markets. They cannot be replicated in the capital markets, and to that extent there will be variability between the assets and the liabilities. That noise, I've learned from my pension experience, meant that you can add equity exposure and get a higher return with very little addition of overall risk. That meant that perhaps insurance companies could maintain much higher equity exposures than they were used to running with.

Also, I found that asymmetry in pension finance is very important, and I believe that there's also asymmetry in the insurance world. By asymmetry I mean that if there is a penalty to surplus dropping below some certain amount or a diminishing return if surplus grows, that is a kind of asymmetry, and that asymmetry should be used advantageously.

I knew a few chief investment officers (CIOs) at insurance companies, and one of them was sufficiently intrigued by the ideas and our modeling capability to allow his company to be used as a guinea pig in a pilot study. The pilot study was sufficiently positively received, and I was encouraged to present it to a wider group of insurance actuaries. I've obviously changed the numbers here to protect the innocent but I still maintain the integrity of the results. We've done a stochastic analysis of assets and liabilities. We first had to capture the liabilities in a way that we could project stochastically, and then we could get into the main part of our study, which is what I will now share with you.

Because of this asymmetry, we also wanted to look at what changes in asset allocation would ensue as the amount of surplus changed. Also, was there such a thing as an optimal amount of surplus? So we started off assuming that we only had assets equal to the statutory reserves. What we used as a proxy for statutory reserves is the discounted scenario, one cash flow at 6%, and that amounted to $17.6 billion. We then looked at what the results would be if we added a further 10% of those liabilities as starting assets, 20–30%.

For every capital market in the area, and we did 400 independent scenarios with stratified sampling to make sure we were covering everything that could happen in the capital markets, we projected the assets and paid off the liability cash flows to see what surplus we would have at the end. We did put in a penalty function. If we ran into negative cash flows, and what we meant by that was if the market value of the assets was less than the statutory reserves at any point of time, then you had to put money in. But we accumulated that at 10.5%, which was 3% higher than the long bond rate, which at that point was 7.5%, and that was our penalty function for running out of money.

Chart 1 is a distribution of ending surplus if you invested 100% of the statutory reserves in 99% fixed income with a duration of 12. This was the distribution we got of the ending surplus at the end of 30 years. The surplus is the market value of assets less the market value of liabilities that were left at the end of that 30 years, less all the accumulated negative cash flows. You can see that it's a distribution that's quite skewed. On average there is a mean surplus, and that's because you actually start off with an economic surplus. The statutory reserve was calculated at 6%. We were starting off in an environment in which the long bond was at 7.5%. But there are many cases where, in fact, we run out of money, where there is a deficit.

CHART 1
DISTRIBUTION OF ENDING SURPLUS
99% FIXED INCOME, DURATION 12, NO ADDITIONAL CAPITAL

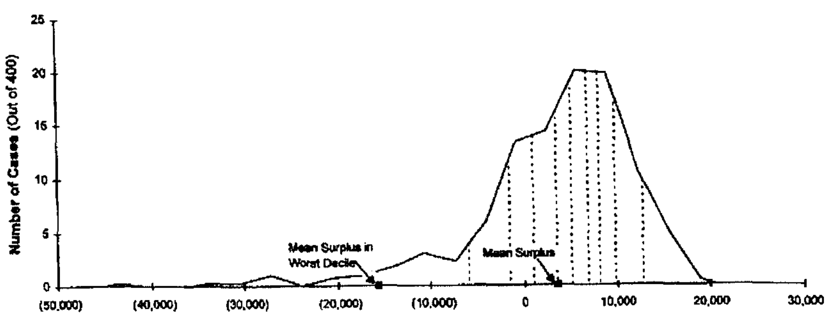

Note: Ending surplus is market value of assets less market value of liabilities less accumulated negative cash-flows.

We're also going to look at many different types of asset allocation, and for that reason, we can't compare these entire distributions each time as it gets far too messy. We need to build this distribution down to a single point so that we can look at efficient frontiers. We took the mean surplus across all the 400 scenarios (Chart 2), and plotted that on the vertical axis, and as a measure of return is the expected surplus. We plotted the mean surplus in the worst decile (in the worst 10% of cases; there were 40 total cases) as the measure of risk on the horizontal axis.

CHART 2
STARTING EFFICIENT FRONTIER
99% FIXED INCOME, DURATION 12, NO ADDITIONAL CAPITAL

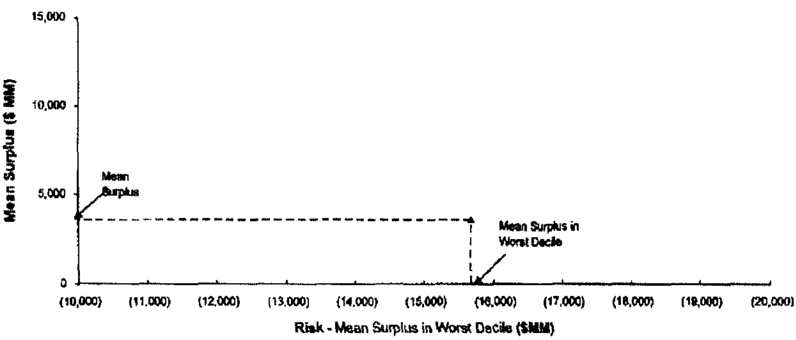

Note: The entire distribution of surplus can be reduced to a single point by plotting the average surplus over all scenarios (return) against the average cost in the worst decile of scenarios (risk). Also, plotting these points for various asset mixes constructs an efficient frontier.

The expected return (mean surplus) is just under about $4 billion (Chart 3). The mean surplus in the worst decile, which of course is a large deficit here, is the risk (Chart 4). It's a measure of how bad things can get. Now we can start plotting an efficient frontier (Chart 5). I'm going to look at what happens to the surplus as we move duration around. I guess our first insight, although it may not be a great insight, is that some of the answers that we got are quite robust. Others are quite soft. I hope you'll find them all interesting and challenging.

Mean surplus is on the vertical axis. That's our proxy for return. You want that to be as high as possible, to get upward on that axis. On the risk axis is the mean surplus or deficit in the worst decile as the risk measure. You'd like to get that to the left as much as possible, and the first thing that comes across is that any duration below the Macauley duration of the liabilities, which in this case was about six, is inefficient.

There is no payoff for having a duration below the Macauley duration of the liabilities. There is, however, a payoff for having a duration above the Macauley duration. There's a risk reward trade-off. The efficient frontier essentially starts at the Macauley duration, but higher durations may still be justified. This is because longer-duration bonds have a higher expected return than shorter duration bonds if you believe, as we do and as most people do, that the mean reversion yield curve is upwardly sloping.

CHART 3
DETERMINING THE DURATION
99% FIXED INCOME, NO ADDITIONAL CAPITAL

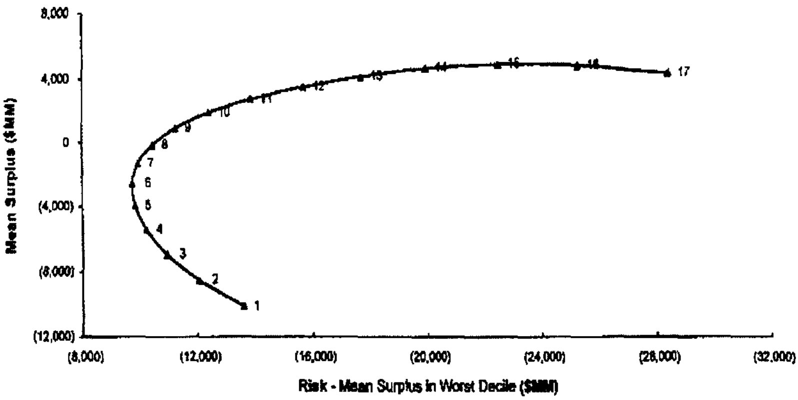

Note: Macauley duration of Scenario 1 cash flows is the lowest of the set of efficient durations. Also, runs with duration of less than 12 run out of money more than 40% of the time.

CHART 4
DETERMINING THE DURATION
99% FIXED INCOME, NO ADDITIONAL CAPITAL

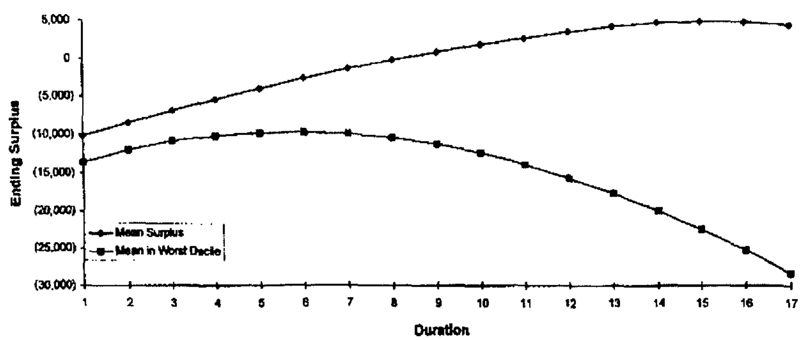

Note: There is a trade-off between minimizing risk and maximizing surplus.

CHART 5
EXPLORING DIFFERENT ASSET MIXES

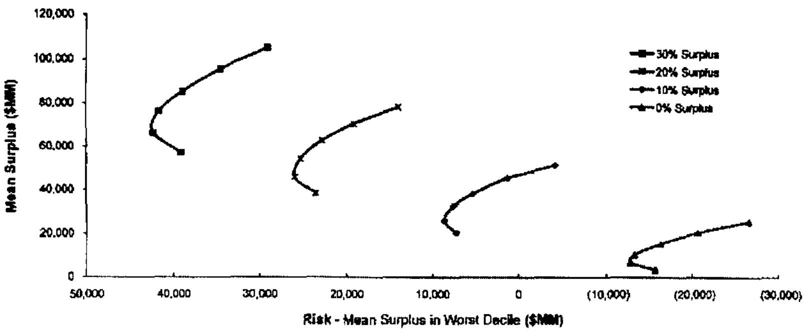

Note: Efficient frontier steepens as surplus increases (that is, justifies more equity exposure). Also, one cannot tell if more surplus is more efficient.

After finding the optimal duration, and we had to actually do this for each surplus measure and each equity exposure, I'm going to look at what happens when you change the fixed income and equity exposure. I'm first just going to focus on that bottom, right-hand graph, the one that says 0% surplus (Chart 6).

CHART 6
EXPLORING DIFFERENT ASSET MIXES
EFFICIENT FRONTIER—0% AND 10% SURPLUS

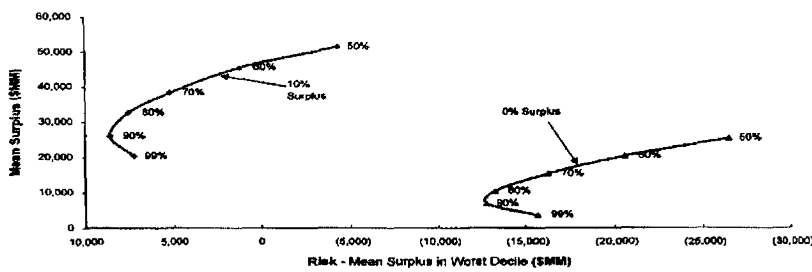

Note: Fewer failures leads to better risk/return trade-off. Also, "noise" in liabilities causes lowest risk portfolio to include 10% equities.

Assets start exactly equal to the statutory reserves there. The first thing that comes across is that 99% fixed income, even though you may not be allowed to have more than that, is inefficient. In fact, the lowest equity exposure that ought to be on the efficient frontier is 10%, and the reason for that is the noise between the assets and liabilities.

If you cannot exactly match your liabilities, if there is noise between the assets and liabilities, and if you add on equity exposure, you get an enhanced return for only a small increase in risk if the equity risk and the risk that you have between the assets and liabilities are not correlated because you get a diversification gain. Furthermore, higher equity exposures are still on the efficient frontier.

Does everybody know how to read this efficient frontier? Should I just go over it in detail? For instance, at 90% fixed income—that's what 90% here means on this bottom, right-hand graph—the expected surplus is just below $10 billion. The mean of the worst decile is about -$12 billion, a $12 billion deficit. That's better than the 99% number, which is just over a $15 billion deficit with a lower expected surplus, and that's why it's dominant. Ninety percent is just better than 99%. You get a higher expected surplus at 80%, but you have to take a little more risk. Your surplus in the worst decile drops from $12 to perhaps $13 billion. But you get quite a big pickup in expected surplus.

It's 10% of your statutory reserves to start off with. If you had a 10% surplus, the efficient frontier changes to the one in the upper, left-hand sector. The numbers all get higher. There are fewer failures. You're getting a better risk return trade-off. The efficient frontier steepens. It might be hard to see, but it does steepen. That means you're getting a better trade-off for more equity exposure, for more risk. Once again, the lowest point on the efficient frontier is 90% fixed income, but higher equity percentages have a better payoff.

We will look at 0% surplus (Chart 7), 10% surplus (Chart 8), and 20% and 30% (Chart 9). The story just continues. The efficient frontier steepens as surplus increases. It raises an interesting question. Can you get a better-than-a-capital-markets-line return ever by adding surplus to any insurance company? Does it ever pay to actually put surplus into an insurance company? Can you get a better return than what the capital markets like? That would happen because of asymmetries in return. If you get either a diminishing return or an increasing loss, this would happen. In our model we have built this in because we have a penalty function for running out of money. In the real world I believe there are these penalty functions. There is an asymmetry in insurance finance, and it leads to some very interesting answers.

CHART 7
RETURN ON 10% ADDITIONAL CAPITAL
0% SURPLUS AND 99% FIXED INCOME

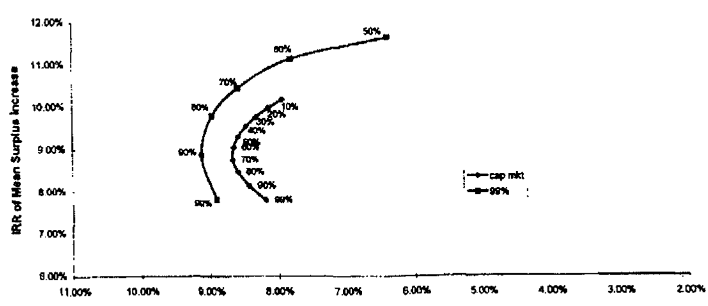

Note: Fewer "borrowing" scenarios, worst decile results are better, so risk returns are high. Also, there is a large payoff for going to efficient asset mixes.

CHART 8
RETURN ON 10% ADDITIONAL CAPITAL
10% SURPLUS AND 90% FIXED INCOME

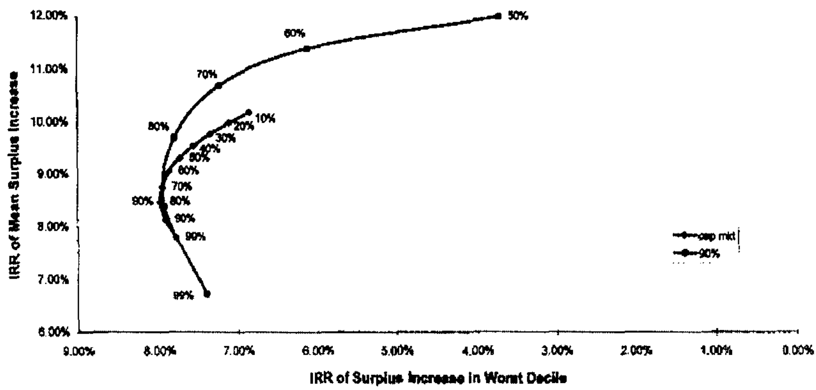

Note: No extra return for increasing surplus without increasing equity exposure.

CHART 9
RETURN ON 10% ADDITIONAL CAPITAL
20% SURPLUS AND 80% FIXED INCOME

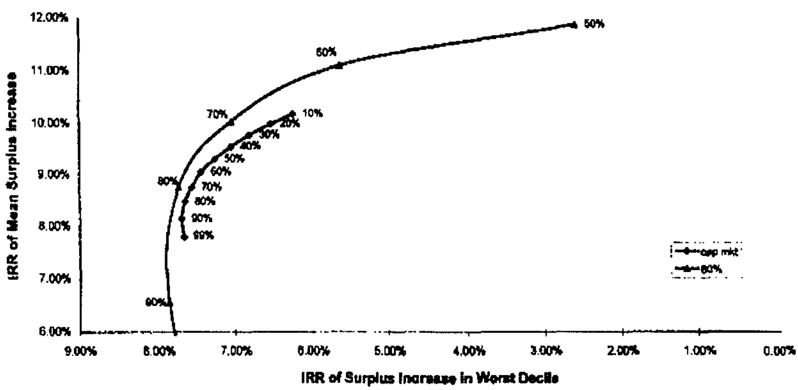

Note: Small payoff is due to "duration noise."

First, if you are going to add surplus, is there such a thing as optimal surplus? The invested capital, to justify the investment, must pay better than investing in the capital markets. We first determined what the internal rate of return is on the same kind of basis if you invest in the capital markets. We then determined the internal rate of return, if you put in the money to cover the insurance liabilities in the insurance companies. We compared that internal rate of return on the ending surplus by increasing the capital versus the internal return if the increase was simply put in the capital markets.

We first looked at 0% surplus with 99% fixed income (Chart 7). We assume that's the practical result. If you do have 0% surplus, you can't have much less than 99% fixed income. If we take 10% of the statutory liabilities, and we invest it just in the capital markets—this is the return line on the right-hand side—the capital markets return, the return if you invested that in 80% fixed income, 60% fixed income, etc., it shows that the expected return is on the vertical axis, the mean return on the amount invested, and the return in the worst decile is shown on the horizontal axis as the measure of risk. You kind of get the efficient frontier that way.

If, however, you put the money into the insurance company, you get a dominant set. For instance, if you now switched the assets to, say, 80% fixed income, 20% equities in total (Chart 8), and you now have 110% of your statutory liabilities as assets, you're getting an internal rate of return of close to 10% on an additional amount that you put in, with the return in the worst decile at about 9%. That's clearly better than putting it in the capital markets, no matter what equity percentage you choose for the capital markets where you cannot get that kind of trade-off. That left-hand line simply dominates the right-hand line. And that would mean that if our penalty function was correct, which of course it may not be, this is a guess at a penalty function, it would pay companies to put money in to be at least at 110%.

In switching from 110% to 120%, there's still a return if you're willing to take more risk, but the return is lower. The extra return you get starts approaching the capital markets line. If, in fact, you go one step further, from 20 to 30, you hardly get any payoff at all. In fact, the only reason why we were showing a very slight payoff is because of a noise factor. The durations were different with the higher surplus than in the lower surplus, and that was not really a hard number.