RECORD, Volume 26, No. 1*

San Diego Spring Meeting
June 22-23, 2000

Session 81PD
Stochastic Immunization

Track: Investment

Moderator: JOSEPHINE ELISABETH MARKS
Panelists: STEVEN LANE CRAIGHEAD
SCOTT E. NAVIN†
Recorder: JOSEPHINE ELISABETH MARKS

Summary: Today's computers allow stochastic modeling techniques to efficiently and effectively consider a wide variety of scenarios and outcomes. How can these tools be used to optimize portfolio performance?

Ms. Josephine Elisabeth Marks: Scott Navin is investment policy officer at John Hancock. He is a team leader for the ALM practice and is responsible for portfolio coordination, ALM research, and oversight of their futures hedging program.

Scott is not a member of the SOA but he is a CFA, so we are pleased to have a colleague from the "real" world of investment professionals here to talk to us.

Mr. Scott E. Navin: My remarks are going to focus on the management of interest rate risk within a market environment. Although we all have to deal with GAAP accounting and statutory reporting, I'm going to concentrate on the process whereby you measure true market risk.

My objectives are to give you a sense of how you can modify the inherent risk in portfolios, review a method of managing interest rate risk, and then add a few comments on the practicalities of dealing with this type of strategy.

I want to begin with defining our sample portfolio. This is a very simple portfolio. We have two assets, a $50 million, 10-year bond and a $25 million, 6-month commercial paper (CP) that's supporting a 5-year bullet GIC. These are all non-stochastic very simple examples; embedded options are not a problem here.

†Mr. Navin, not a member of the sponsoring organizations, is Investment Policy Officer at John Hancock Financial Service in Boston, MA.

Note: The chart(s) referred to in the text can be found in the documents 081sd1pd charts 1-4, 5-8, and 9-10.

Chart 1 shows us the portfolio cash flows associated with this portfolio. You can see we have short assets and long assets balancing our GIC. This is a duration-matched account. It's not cash-managed.

Chart 2 shows us the yield curve. We've got three points, 0.5-year, 5-year, and 10-year; the rates are 6%, 6.8%, and 7.2%. Even though our example is very simple, the method we will talk about is independent of the yield curve points and it's also independent of what kind of interpolation you've made between points.

Now that we have a portfolio and yield curve, what we need is a technique to start the valuation process, and I want to describe the most common approach, and this involves a parallel shift approach where the curve shift is parallel, up or down.

The price is a function of the change in the interest rate, and modified duration; in this case it's the first derivative of the price function. The reason this is important, obviously, is that derivation is at the heart of interest rate risk management.

A lot of this is probably familiar to you, but I hope you'll bear with me. In terms of approximating duration, what is the calculation? We can use the central difference method; whereby you shock the curve up and down a few basis points. Then you value your securities. This is a good technique to use, because if you do have embedded options, you typically can't produce good cash flows. Stochastic cash flows are next to impossible to obtain and deterministic cash flows are usually worthless. This technique avoids the issue of having to deal with cash flows.

Now we have a portfolio, our yield curve, and our evaluation technique, let's look at some of the portfolios and statistics associated with this simple portfolio.

Table 1 shows the asset liability management balance sheet. You can see the market value of the portfolio assets and liabilities for surplus of $9.29. The asset and liability durations are approximately equal. If they were equal, the surplus would have the same duration as the assets and liabilities; it's actually a little bit longer because the assets are a touch longer than liabilities. There's a lot of convexity.

TABLE 1
PORTFOLIO VALUATION
(On initial curve: 0.060,0.068,0.072)

	Market Value	Duration	Convexity
Assets	$80.68	4.93	21.37
Liabilities	-71.39	4.88	13.03
Surplus	$9.29	5.31	85.42
Surplus Ratio (SR)	11.52%		

You'll also note that the surplus ratio, surplus relative to assets, is 11.52%. Typically, if you wanted to have a surplus ratio that is impervious to interest-rate movement,

you would want the assets and liabilities to be equal. If you're looking for classical immunization, which would be a surplus of zero, you have to align the duration of the assets to be consistent.

Duration allows you to fairly accurately estimate market value changes on your portfolio for a given change in interest rates. It's a good technique for estimating the effect of small parallel changes in the yield, which unfortunately, is not typically reflective of reality.

But in any case, in our example, the surplus is $9.29, with duration of 5.31 and convexity of 85, a positive 100 basis point movement would reduce the surplus to $8.84.

Given some non-parallel yield curve shifts, one to steepen and one to invert, let's take a look at what our new portfolio looks like once we value it (Table 2).

TABLE 2
SURPLUS RATIO VOLATILITY

	Initial	Parallel	Steepen	Inverted
Assets	$80.69	$82.72	$77.03	$84.42
Liabilities	-71.39	-73.16	-69.55	-73.49
Surplus	$9.29	$9.56	$7.48	$10.94
Surplus Ratio	11.52%	11.56%	9.7%	13.0%

In the steepened environment, our surplus ratio changed fairly significantly to 9.7%, and in the inversion, it goes to 13%. There's quite a bit of volatility, even though the curve isn't doing a whole lot. What you'll notice, though, is that your surplus reduced from 11.5% to 9.7%, implying a 300 basis point shift. There's some leverage there that catches you if you're not careful. You need the ability to manage the non-parallel movements.

Essentially, as we just discussed, the surface ratio is not immunized, and it's going to bounce around. The predictive ability of duration is not at all what it seems when you have non-parallel moves, which is the standard world.

The questions become, how do you quantify this non-parallel risk, and how do you try to manage it? To do that you need a slightly more advanced valuation model. A non-parallel shift approach essentially builds on the duration approximation. You use the yield curve as a model, with independent or multiple shifts along the yield curve at key points, traded points typically.

The price function of the model is a direct result of the rate changes at particular points. The big difference is that instead of having 1 duration, you have partial durations. The partial durations allow you to focus at particular points along the yield

curve, and see the true sensitivity of the portfolio to the various moves along that yield curve.

How do we actually obtain these partial portfolio durations? Essentially, if you're interested in the five-year point, you shock the yield curve up and down at the five-year point, and in each case you use your spot curve or zero coupon curve, do your evaluation, and work out the duration.

Table 3 shows us the partial durations for our portfolio at 6 months, 5 years, and 10 years. The five-year partial duration is –38 while the 10-year partial duration is +40. If you have a non-parallel move, you can see why things get off significantly. As you can see, the total partials add up to your original portfolio duration.

TABLE 3
PORTFOLIO PARTIAL DURATIONS
(Initial curve with 5 basis point rate shocks)

	Assets	Liabilities	Surplus	Surplus Ratio
D_1	0.18	-0.32	4.03	0.50
D_2	0.16	5.20	-38.56	-5.04
D_3	4.59	0.00	39.84	4.59
Total	4.93	4.88	5.31	0.05

Now we have partial durations. What can we do with them? What good are they? The beauty is that you are no longer restricted to that artificial rule of parallel move in the yield curve. Now you can accommodate any move in the yield curve. That will allow you to extract a lot of information out of the portfolio. We can get a sense, given this partial-duration profile, of how it would move, over historical periods of time.

Once you have that kind of information, you can get a good sense of how you might want to build this into your investment policy. What kinds of constraints are appropriate, how to establish risk return tradeoffs, to the extent that everybody is always looking for more yield. We always want to go longer if possible, in an upward-sloping environment.

But there are consequences. This allows you to say, "Okay, if I extend another half-year, what is the implication in terms of the standard deviation of the portfolio, versus what I expect to pick up using the current income?" You start making trade-offs.

What you see is that you can fairly readily predict the impact on the pricing of your portfolio. For instance a non-parallel shift of –2 bp at 0.5 years, +19 bp at 5 years and –20 bp at 10 years increases the market value of our surplus to $10.72.

It might be interesting to look at the equivalent parallel shift of that small move in the yield curve. You can do that in your head. Those small moves equate to having a

290-basis point parallel move, which is quite amazing and wouldn't be intuitive to most people, myself included!

How large can that equivalent parallel shift become? I'm going to use the concept of durational leverage, which says that based on a vector change in the yield curve, and a vector of partial durations, the maximum limit is defined.

The formula below shows us the derivation of durational leverage (L) for this example. How does that reflect reality? We ran this portfolio against 23 years of actual data, quarterly overlapping data. We calculated the rate movements at the particular points along the simple yield curve, and then worked out what the implied rate movements were.

$$-L|(i,j,k)| \le F \le L|(i,j,k)|$$

$$|(i,j,k)| = \text{length of vector} = [i^2+j^2+k^2]^{0.5}$$

$$\text{General: } L = |(D_1,D_2,D_3)|/D$$

$$\text{Example: } L = |(4.03,-38.56,39.84)|/5.31$$

$$=55.6/5.31$$

$$=10.47$$

$$-10.47|(i,j,k)| \le F \le 10.47|(i,j,k)|$$

Chart 3 "Equivalent Parallel Shifts" shows the results of this. It's quite interesting. It turns out that the parallel shifts approach 300 basis points in several quarters. Typically, the standard deviation is more like 100 basis points. But you do get some significant moves.

In a similar vein, it would be interesting to see what this portfolio would have done in terms of its standard deviation over that same period of time. To do that, you need the Treasury database, you need to plot the changes at particular points, and then you apply that back to a simple formula.

Chart 4 shows the actual distribution of surplus changes that would have occurred in our portfolio for this 23-year period.

Again, they look fairly normally distributed, with a little bit of a tail, but in any case, the standard deviation for this portfolio is about $0.5 million, which is 5% surplus. The shift that we discussed earlier, based on a shift of minus 2, plus 19, minus 20, was a little over 15%. As you can see, it's way out in the tail.

If you wanted to start managing to it, you could use a closed formula solution. Once you've calculated all the data points, you can create a co-variance matrix; we called it "K," giving monthly, quarterly, and annual rate changes for your decision process. Calculate that, and then you've got fairly simple matrix multiplication, to get your standard deviation.

What it's saying is you take your market value of surplus, your $9 million, and you multiply that by a row vector of partial durations, times the matrix, times a column vector of the partial durations. You don't take the square, that's your variance.

When you do that, you get a sense of what the surplus volatility is, in that particular portfolio. Again, if you're comfortable with it, if it fits your investment guidelines, or your investment policy, then you're in good shape. If not, you need to manage it.

What you can do is a little more in the way of stochastic immunization. Here what you want to do is minimize this risk, subject to constraints that are important to you. For example, you might want to go to senior management and say that for this portfolio, there won't be an impairment, or loss of surplus greater than 5%, 95% of the time. The way you can get there is on that previous formula, that variance that we just described, you can minimize that with constraining optimization.

Basically you can layer on whatever constraints you want. Typically, we would be concerned with the size of the trade that we are contemplating.

In terms of a trade, we are usually talking about swaps or options, something of that nature. you would have an overlay on the portfolio, so you don't disrupt the portfolio management that is in place.

You would usually want your strategy to be cash neutral, and if it's a duration-managed account, you want it to be duration neutral as well. You might also want to have notional neutrality because you don't want to have huge current income implications.

To consider at a simple example, you want to minimize the variances in the standard deviation, with no constraints. You have 1, 5, and 10-year swaps that are available. The optimal swap trade is to pay fixed on $43 million notional at 1 year, pay floating on $83 million at 5 years, and pay fixed on $48 million at 10 years. Your total notional amount is $170 million, and you probably wouldn't want to do it.

But if you layer that swap on the portfolio, you end up with partial durations that are zero. If you get there, you are basically bulletproof, because the curve can do whatever it wants, and it's going to have no impact on you.
If you don't want to do the whole trade, if you do a portion of it, 25–50%, you can bring the tails of the standard deviation in, which is helpful. When you actually do this kind of work, you'll find that you can accomplish a lot with the initial trades. After that it's a diminishing return.

A practical consideration is that this is only as good as your cash-flow data. You want to make sure that your valuation processes are all the same, so that if you're getting data from different places, they are all being done on the same format, and that the option models are consistent.

Counter-party exposure is there, because you're going to be layering on derivatives. Once you start going down that route, unless you are using exchange-traded futures, then you've got counter-party exposure. The fact is you are basically trading off interest rate risk for credit risk. You need to be able to manage that, or quantify it, first.

For FASB 133 compliance, you've got to take those macro portfolio hedges and convert them into specific micro hedges. You need some kind of algorithm or process that will allow you to do that. Otherwise, you are going to introduce huge GAAP income volatility into the process.

Finally, you have to manage expectations. You've got to be wary of senior management thinking that you've got everything covered. Because even if you do a good job, and you think there's only a 1 in 100 or 1 in 1,000 chance of failure, those things seem to happen more frequently than that, somehow.

Ms. Marks: Steve Craighead is an ALM actuary at Nationwide. He has published papers, and is a regular speaker on optimization, applied chaos theory, and various other topics. Steve is currently a member of the Investment Section Council, and has been a fairly regular speaker at SOA events over the years.

Mr. Steven Lane Craighead: I'm going to present a concept that we've been using at Nationwide and probably other companies have been doing it too. I'm going to take on stochastic pricing as a possible alternative to immunization. If immunization is used to back a line of business, there's a whole series of assumptions to be made.

- One can calculate the duration or convexity of liabilities. But how do you discount it? Dave Becker has written about these issues in the *Transactions*.
- Liability cash flows are usually pre-tax and usually pre-risk based capital.
- There may be little or no consideration for the schedule of earnings.
- How do you value the portfolio after the first period?
- What about sensitivity analysis?

We should consider actuaries to be artists not scientists, because we are working with psychological issues such as policyholder distribution channel management. That creates a whole realm of problems.

Stochastic pricing might address the above issues, even though it is more complicated. It's like the Anderson pricing method that people use to price their products, and you usually end up having one scenario in your price based off that, and then you're suddenly taking assets, liabilities, and surplus into consideration.

Well, in some sense, taking that approach is more complicated. It does allow one to consider various product revisions as well as portfolio selections. You could change some option on a liability model, or you could actually change the portfolio selection

that you are putting into your asset model.

We found that some of these concepts have been very effective to describe stochastic results to upper management. If you are really excited, you may get upper management to think in terms of standard deviation, but you rarely can. However we've got some ways that we've been able to housebreak them a little bit.

Your target parameter to optimize could be the present value of distributable earnings (PVDE) or the present value of accumulated surplus. I'll just refer to the present value.

Your audience may be the regulators, who are concerned with your solvency margin. You could use it to manage your policyholder guarantees, or you could use it to manage your shareholders' expectations for dividends.

In my company the upper management might like a nicer return, because their bonus improves. You can take into consideration all these different levels of whatever you are analyzing, and set your target.

Probably most of you have seen the "S" curves where you take your different present values of whatever you define as your target, and sort it from lowest to highest. Chart 5 shows us the "S" curve for the PVDE and from it you can assess the probability of not meeting your returns.

Chart 6 shows us the distributable earnings through time. If you are concerned about your earnings volatility, you might want to look at the cash flow through time, not just the present values.

Chart 7 shows the distributable earnings for all scenarios at a specific time "t". It's a vertical slice of the distributable cash flow at that time. You can see the behavior of your model through time and whether you have earnings stability. The way that we've found to be effective to explain these types of models to upper management has been to do overlays of these different graphics.

Chart 8 compares the pattern of distributable earnings for two projects. Project One may not be quite what upper management wants, because Project Two has better performance. They might live with that little bump. But sometimes this method has been very effective at convincing management so they can understand it.

You can also look at the "S" curves in Charts 9 and 10 for these two projects. The "S" curve would definitely say to take Project two, because in all situations, it has done better. The "FISH" diagram shows a narrower range with fewer losses for Project Two.

Remember you can consider the needs of the regulator, policyholder, stockholder, and upper management, you could do the same things with these kinds of diagrams.

Instead of slicing it at zero, you could slice it at higher levels based on whom you're trying to please.

But I want to use an example of all this to show you something that we did, to look at a choice of a portfolio for Bank Owned Life Insurance (BOLI). I may skip some of the details, but if you want more detail you can go to our Web site at www.a-l-m.com.

What we're going to do is find the static portfolio strategy, where the money is going to be initially invested at that strategy, and on rollover. You'll also invest at the same strategy. Then with the dynamic portfolio strategy, we're going to determine what we're going to do if an economic variable goes above some limit. If it stays within the limit, we'll use one static strategy, but if it pops up, we will use a different one.

What we're going to determine in a dynamic strategy, is a portfolio mix below the limit, the portfolio mix above the limit, and where the limit should be set. It's a fairly simple dynamic strategy. The further details are in the paper, "Portfolio Optimization in Corporate Models" by W. Babcok and S. Craighead (November 1998).

We've made some simplifying assumptions:

- The portfolio selection only uses the efficient frontier.
- The interest rate scenario produces realistic scenarios (using Mark Tenney's DMRP™ model).
- Expense and decrement assumptions are all correct.
- The asset universe is limited to non-callable bonds with static spreads over Treasuries.

Also, we did not modify the liability model. We just changed the portfolio allocation. However, we did look at the statutory distributable earnings.

We needed to set up an optimization rule to maximize the internal equity and minimize the downside present value of distributable earnings. We realized that this reflects only our personal utility preferences, and not that of our management.

There is a problem. This does take a long time to run. However, there are some developments afoot that can speed this process up. Then you can determine the portfolio structure for new cash flows.

I'm also doing research on how to use a restrictive linear model so you don't have to go through a million corporate model runs. The risk was defined as being the present value of distributable earnings falling below zero. In other words, I'm looking for downside rules, and I don't want to affect my upside potential. My downside risk measure looks at lower second partial moments (partial standard deviations) defined as:

$$\text{Parstd} = \left\{ \left(\sum \min\left(x - \text{ave}(x), 0\right)^2 \right) \Big/ (N-1) \right\}^{1/2}$$

One problem that was encountered was that optimization would occasionally create a portfolio that would short assets.

There is some justification for allowing negative asset allocation. You could maximize the economic value of the line of business. Negative allocation may be counterbalancing the underlying options in the liability. The revenues can be used to create an equivalency to shorting; that might be allowable.

If you want to eliminate the negative allocation, just make sure the other allocations are reallocated so that you have a total 100% allocation. Your universe gets smaller, and your run times get faster.

However, if you want to retain all your assets in the universe, and you can't have negative allocations, you've got to use some transformation of the optimal asset allocation so the percentages are always positive. Use log normal.

Whenever you do optimization, you've got to come up with a target function. I want to pick a target function between zero and one. My target function is going to consider both return on equity and partial standard deviations and I want to put twice the emphasis on return on equity. My target function is:

$$(1 + \text{ROE})^2 * (1 + \text{PARSTD})^{1/3} - 1$$

If I maximize this function, I maximize the ROE while minimizing the standard deviation.

Now for the dynamic allocation I told you about. We created static portfolios and we plotted return versus risk for each of these to develop an efficient frontier. Using the efficient frontiers, we found that a 3–10 year barbell gave us the best results.

In our dynamic model, we didn't get quite as good results. It may have needed more scenarios. We could possibly make an initial recommendation to a portfolio manager of the type of asset to buy (although our investment people don't like us to tell them what to buy). But we could hint.

It was still an interesting observation. I'm going to paraphrase Bob Reitano, "model building is like making sausage. You enjoy the results, as long as you don't know the ingredients, and how it was made."

Mr. Frederick W. Slater: Scott, you are essentially assuming an economic framework in your immunization, you looked at the difference between the market value of assets and the market value of present liabilities. If you stick that into an actual framework, a statutory framework, I'm not sure whether this will still work, because in the statutory framework you have book-value accounting. The reserves don't change if interest rates change, they are pretty much static. Can you comment on that?

Mr. Navin: Yes, it's an interesting question. All I can say is, we've been doing it for a long time, and you will find it in our annual statements. I think what you need to do is be cognizant of what you are doing on a statutory basis, and build constraints that won't allow you to do irreparable harm to a portfolio. You have to think it through. We don't end up with radical policies. You wouldn't look at it and think that it wouldn't make sense.

Mr. Slater: Anything in particular that you think will drive the problem?

I guess you'll always be fine, because you can sell if you need to match a liability cash flow and still be in a situation where your market value of assets is greater than your market value of liabilities.

Going forward through time, the strategy requires rebalancing. I think that's one aspect that presents a problem, rebalancing in a statutory framework, you're selling at market value, but you're getting a book value.

What makes this even worse is that you have to pay taxes on capital gains. If you try to model that going forward, I could see a lot of problems with that approach.

Mr. Navin: There is an issue of cash flow mismatch. That's one of the real concerns but there are ways to manage it. Essentially, we end up with some of our highly immunized portfolios with a lot of floaters that give us the liquidity. They might be synthetic floaters, but at least they give us the liquidity we need to be able to fund our liability payoffs as we go forward.

Mr. Claus S. Metzner: Steve, my question is whether you've also run this process for multiple legal entities within the group, to ascertain whether or not you can structure an investment portfolio for a group like this, and then shorting assets statutorily. Is it practical or feasible; and if so, what was the management title issue, so that you could get an agreement from the various profit-responsible managers?

Mr. Craighead: We have situations where we have lines of business that lapse, which had huge amounts of assets still associated with that specific group. We

created something, internally, like a negative asset, or a theoretical asset, that everybody could buy or sell between the different lines of businesses.

We thought that might be a good plan, and in many ways it has been. But there is some difficulty in that now, as we change more of our line of business mixes, it has created a surplus problem because some people who have those negative assets are actually borrowing from surplus at different rates of return.

It's a good idea, in some sense, but it's really hard to unwind. I understand what you're saying. You're trying to create an optimal corporate portfolio. But an internal swap has some difficulties, too, because you may have other requirements on specific lines of business, and they may have to have a certain rate of return or a certain guarantee to be funded.

Mr. Metzner: When we deal with these issues there's obviously no single solution. Very often, we are finding, as you mentioned, the unwinding can be problematic, because then you may tend to have a false profitability indicator for new lines of business of the products.

Mr. Jose D. Siberon: For international countries like Japan or Korea where the asset universe is not as good as in the U.S., what kind of immunization technique do you recommend? Especially for life insurance, where the durations are so high?

Mr. Navin: My first thought is that you would move into more efficient markets, and use currency swaps, so you don't need to manage that risk. The reverse of that has been true in recent years. Domestically, our investment departments have been moving outside the U.S. because there are less efficient markets out there; they are obtaining what they believe is good yield, using their credit expertise.

But by the same token, they don't have to concern themselves with currency risks, because we don't handle currency risk in our company. We eliminate it at the issuance of the bond, and it really opens up the universe, and allows you to leverage your internal credit skills.

Mr. Craighead: I have a comment on this. Back in 1995, I went over to Moscow and taught for about two weeks. I had a question very similar to yours. How could we create our portfolios to protect our insurance companies? The greatest difficulty within Russia at that time was that they were restricted to rubles for their reserves. Now the ruble was devaluing at the time, and we didn't have a yield curve past three months. Now you have to look for potential investments that will continue to inflate at the same rate of inflation. My answer in Moscow was that you had to buy real estate near the city.

Mr. Navin: I think you're probably talking about the emergence of asset versus liability spread on a statutory basis. That is an issue to the extent that you're issuing the products. Say you're issuing a 5-year GIC and buying a 10-year asset. You're

implicitly willing to take on corporate spread risk. As the 5-year GIC runs off, you're left with what was originally a 10-year asset now being a 5-year asset. You now have an asset that was issued in potentially a different spread environment.

What happens is, over time, we are in markets continuously, so we have to average things out and let things take care of them. We also make a real effort to make sure that our investment curve is driving our liability pricing, so there is a real tight tie there. We found that the emergence of this asset liability spread on a statutory basis has been pretty stable. The markets are improving all the time.