RECORD, Volume 23, No. 3*

Washington Annual Meeting
October 26–29, 1997

Session 151PD
Asset/Liability Modeling and Portfolio Optimization

Track: Investment
Key words: Investment Products, Investments

Moderator: CATHERINE E. EHRLICH
Panelist: LAWRENCE A. BERGER†
Recorder: CATHERINE E. EHRLICH

Summary: This panel presents an overview of present methodologies and design considerations for asset/liability modeling. The panelists discuss how to identify optimal strategies for risk management.

Ms. Catherine E. Ehrlich: I am the entire actuarial department at Capital Management Sciences. We are a software vendor and we specialize in fixed-income portfolio management software. Many of our clients are insurance companies, but we also work with banks, total return managers, mutual funds, plan sponsors, and so on. I'm going to do the first part of this presentation and talk about a generalized asset portfolio optimizer. Along the way I will hit on a few asset modeling considerations, but if I try to give you an overview of the present methodologies of modeling assets, we'd be here until the next Society meeting.

Larry Berger is going to give us a more detailed presentation on a specialized asset/liability model that he works with. Larry's an economist with Swiss Re New Markets, and he previously was with Milliman and Robertson (M&R).

I'll start off by talking about the different kinds of optimization that are available and use that as an introduction to what optimizers can do for you.

I'm going to talk about one type of optimizer in particular which is a total return optimizer. I'll talk about what specific applications you can use with that optimizer.

---

†Mr. Berger, not a member of the sponsoring organizations, is an Economist at Swiss Re New Markets in New York City, NY.

The software I work with is called BondEdge, so we'll take a look at an example of the BondEdge total return optimizer. On the way, we'll look at some of the constraints that we use with that optimizer and talk about what you do once you have a solution from an optimizer and the impact of the practical constraints as well as the impact of economic uncertainty on results.

What is an optimizer? My own definition of this term is, it's the solution tool for finding the best portfolio for a certain purpose, given a certain set of constraints. I think I've qualified and caveated that enough to really make it almost anything in the world. What you're doing with an optimizer is starting off with a universe of securities, and you, as the user, will define what that universe is. My optimizer is designed to work on domestic fixed-income securities, but, in theory, you could have an optimizer look at any group of securities, even small ones. You can have it just look at this mortgage market or just look at the Fannie Maes in the mortgage market. It can look at something as broad as a wide variety of security types.

The solution will always be subject to user-defined constraints and user-defined constraints can be absolutely nothing. You might not want to have any constraints. You might just want to go out there and let it pick the best thing for you, and then decide if you need to impose some constraints on that. It can be a very detailed list of constraints that you've gotten from your investment management committee or some outside influence, like a rating agency or a regulator, has forced you to have by telling you that you can't do something. It is heavily dependent on asset modeling assumptions. Obviously, you can't really figure out what the best thing to do with your securities is if you don't really know how those securities behave. Everything I'm going to talk about really assumes that you have a way of modeling all securities in the universe, which means if you have collateralized mortgage obligations (CMOs), you're modeling the entire deal. If you have anything mortgage related, you have a dynamic prepayment model. If you have anything with an option involved, you do have an option model available to you. To properly value the mortgage-backed securities, you're probably going to want to have some Monte Carlo analytics available to use. Part of the optimizer has to be all of this state-of-the-art asset modeling.

Finally, you need to come up with some economic assumptions as part of what you're getting at. You might want to sample every single thing that can possibly happen in the universe and come up with one investment philosophy that's going to last forever. You might want to see within certain constraints, what the best way to structure the portfolio is for the next month. You can really take a long-term view or more of a shorter-term view.

Let's discuss different kinds of optimization. The first kind of optimization is yield optimization. If I believe that interest rates are never going to change, then what I want to do is pick the securities that have the highest yield, given certain constraints. I want to make sure that I have the right duration, I want to have the right set of cash flows, or I want to have my securities set up in certain sector buckets, according to whatever my investment committee told me to do. This type of optimization works fairly well if interest rates don't change, which we know doesn't really happen. It works pretty well if there aren't any embedded options in a security. These days you just don't find a lot of those.

A long time ago, my company came up with a yield optimizer that we use for both total return portfolios and for matching against a set of cash flows. We got real fancy, and we put in something that we call the yield convexity rheostat. A dimmer switch is a common example of rheostat. If you wanted to maximize yield or maximize convexity or have some variation of the two you could use this rheostat. That's about as close as we got to really analyzing the embedded options in those securities. We quickly found that as the world changed, and as securities got more complex, people want to use this optimizer, but it became more complex, so it just wasn't being used. We really needed to go and look at new kinds of optimizers.

The remaining two kinds of optimizers are a total return optimizer and mean variance optimizations. A total return optimizer really looks to maximize the horizon return over a given set of scenarios, and that can be the stochastic scenarios that model the entire set of economic assumptions, or it can be shorter. For example, you might say you think interest rates are going to either go up by 100 basis points or stay the same. What it's going to do is look at every single security in the universe, and try to maximize your return, given the set of economic assumptions you gave it.

A mean variance optimization is more academic. What it is trying to do is come up with that efficient portfolio relative to a certain investment horizon. It's generally looking at some ratio of what your expected return is versus what your volatility is. Many of the ones I've seen will define volatility as the standard deviation of returns. I think it has some very good applications if you have a lot of disparate sources of risk for which you can't easily figure out the sources of risk for your portfolio. In contrast if you're doing something like domestic fixed income where you can tell with certainty where your sources of risk are coming from, (90% of it is what is going on with the overall level of Treasury rates), it can really abstract the results away from what the sources of risk are. It makes it harder for me to understand intuitively why the optimizer picked what it did. In that situation you can use duration as a very good way to control your risk or to decide how much risk you're going to take. At times it can be unstable because it depends on correlations

between risks in different marketplaces. Those don't tend to be stable over time. They can be unstable. For the specific application that I'm talking about, with domestic fixed income, we found that the total return optimization works very well and really has a lot of applications.

The uses of total return optimization. We've just recently released this, and I've already found that my clients are far more creative than I am in coming up with uses of a total return optimizer. The first and foremost way I use the optimizer for my clients is to help them formulate a portfolio strategy. You put the constraints in, and the optimizer tells you what the appropriate securities are and what your expected return is. It lets you see, given the constraints you have, where you want to be in the marketplace. It really takes your view of the marketplace and what's going on out there and turns it into specific areas that you want to be in. I've also had clients use this once they decide what their strategy is to kind of see how much it's costing them to have certain constraints. In other words a client might say, "I always want to be 25% in Treasuries; that's the minimum amount I want in Treasuries," and the client's investment people might say, "That's really costing a lot of money." You can run optimization with and without that constraint to see what the difference in expected return is and really see what the constraint costs you. It might be fixing a problem for you but it might not be worth it for the amount of return that you're giving up.

The next use of the optimizer is what I call identifying attractive opportunities. If you have an inventory list or some list of securities that you're interested in pursuing, you can put that into an optimizer as your universe and see within that list where the opportunities are. You can say, I know I want to have some segment of my portfolio in the mortgage-backed market, but I'm just not really sure where the opportunities are. You can really limit the optimizer to just looking at opportunities within that market. Given your set of economic scenarios, it can come up with a suggested portfolio for you. You could have a much more specialized application.

The final use is probably the one that gets used the least. I call it constructing portfolios in detail. If you had a new inflow of cash, and you want to decide exactly what securities to buy, you can give it a list, and it can come up with exactly what you can buy. You need to use that list with a certain amount of caution. Optimizers are very good at finding mispriced securities. If you have a problem with some data in your universe, it will find that because it will probably have a very attractive yield or the security will behave very attractively over these economic scenarios. So, you do have to be careful. We use this most often to get an idea or a strategy, rather than to actually say these are the securities that I know I'm going to buy. As a vendor of software and not a seller of bonds, I can't tell

people I know what bonds are out there, but we do give them ideas on where to go in the marketplace.

When you use an optimizer, I think of it as a process rather than something to use to come up with a solution and be done with it. It's not really a mechanical process but it does require some judgment. Start off by defining your bond universe by deciding what you want to look at? Do you just want to look at Treasuries? Do you just want to look at Treasuries and agencies? Do you want to look at a wide spectrum of securities, or do you want to really limit what you are looking at to maybe rebalance your current portfolio and just use that as your universe of securities?

The next thing you need to do after that is define what I call your hard constraints, which are the things that you don't have the authority to change. For instance, you may have to say, "I can't assume that the optimizer will let me sell off my private placements. That's not a realistic thing to do." You need to constrain the optimizer from selling certain things. I may have certain constraints that are passed down to me by rating agencies or by regulators who say I really can't get into this part of the market because it's going to lower my ratings. There are certain constraints that you just don't have the authority to change or the desire to change, and that's what I call your hard constraints.

Next you want to set up your flexible constraints. Maybe you have some sort of strategy that you've been using overall. Maybe you want to be 50% in corporates and 25% in mortgage-backeds, or maybe your duration should be somewhere between five years and six years. You really aren't set up to know exactly where you are going to be within these constraints. I've set those up next to my flexible constraints.

Then you need to set up your market scenarios. Generally, when I do an optimization, I'm not stochastically sampling the universe. I'm not trying to find the ultimate strategy that's going to last forever and ever. I'm doing more short-term optimizations. All my clients will come in and say, "I think that we have a trading range of interest rates. They're going to move up 50 or down 50, but I'm not really thinking about these other real extreme scenarios. I want to find something that's going to work for me for the next six months." I come up with some market scenarios and in all probability weight those scenarios. It's not just an overall level of rates that are changing. I can define the overall level of rates to change, but I can also change the shape of the yield curve. I can change sector spreads, mortgage spreads, and a number of things in the marketplace.

The next thing I'll do is run the optimizer. These first four steps give me my inputs.

The BondEdge optimizer runs through and might pull out really good bonds for you to buy. I will analyze the results because I want to find out if I have a mispriced bond in there. If BondEdge is telling me to buy 100% of some bond because it has an amazing yield and it's never going to have a loss no matter what happens to your economic scenarios, then I probably have a problem modeling that security or it might tell me to buy things that I believe are rather illiquid, and I don't want it to buy those. If so, then I go back and I take those out of my universe. Generally we talk about going back and adjusting your flexible constraints and adjusting your market scenarios. The first step is to go back and adjust the bond universe to see what garbage I've let creep in. Once I look at the results I can see where I'm hitting up against constraints and I might wonder why I put a certain constraint in because the optimizer might want to go beyond that. I'm hitting up against the duration constraint or a sector constraint or some maturity band I've put in, and BondEdge thinks it can do a better job if it really pushes through that constraint.

I might want to adjust that constraint a little bit to see how I can improve my return by just playing with the constraints. The optimizer is a very objective tool used to look at all these inputs that you've made. You may, when you see the results, think, "I didn't really think I was biasing my view of the marketplace or what was going to happen with interest rates, but based on what Bond Edge picked or what the optimizer picked, I can tell that I really have a very biased view of the marketplace, and I wasn't intending to put that in there." So, you may want to go back and fiddle with your market scenarios as well.

My first warning here is not to treat the optimizer as a black box. You really limit utility of the optimizer if you let it figure out the exact bonds that you're going to buy. The more interesting and the more effective way of using an optimizer is to try to understand what it was trying to tell you about what's out there and about your expectations of the marketplace. I always try to figure out what the optimizer is thinking. What is it trying to do? Why did it come up with the suggested optimal portfolio? What is it about this portfolio that's better than something else? That can often lead you to clues on what you really want to be doing with your portfolio rather than trying to limit yourself to the buy list groups, saying, " I can't buy this bond. I'm really in trouble. I don't know what to do." I often compare the original portfolio I started with to what I ended up with to see how it's different. What is it asking me to sell? What is it asking me to buy? Why would it think that this security is better than the other one?

I'll also use a range of detailed portfolio risk measures to kind of help me figure out if I've brought more risk in the solution than I had really intended to do. I'll look at key rate durations or something to help me assess the risk of slope changes, if I hadn't really intended to look at that in the optimizer. I may look at a vega

measure. What happens if my volatility changes? How will that affect the price of my securities or the value of my portfolio? I may look at a prepayment uncertainty number which will tell me how sensitive my portfolio is to a change in a prepayment model, because it is just a model. What if everybody starts prepaying differently than they have in the past? What does that really do? If the optimizer loaded me up on mortgage-backed securities, maybe I'm taking on a lot of prepayment model risk that I hadn't intended to do. I can look at a spread duration number, which tells me how sensitive my portfolio is to a change in market spreads. All these numbers can help me get a better handle on what the optimizer has asked me to do. Keep trying to talk about it in a general way, but 100% of my time is spent on BondEdge. Any optimizer can do this for you.

Optimization doesn't replace your judgment. I have a lot of clients that want to buy the optimizer because they think they can just load all these scenarios in, and it'll tell them what to buy, and they won't really have to worry about it. It really doesn't replace your judgment. It'll give you ideas, which is why it's often very popular with brokers who are always trying to find ways to help you trade your portfolio. They are always looking to run things through an optimizer and say, "You can get a little bit more yield if you do this." You can also take the results and look at them through other tools. For instance, I can take the suggested portfolio, and I can run it through my cash-flow testing model. I could put it through a value-at-risk test. There are a lot of different tools we have at our disposal to help us analyze how much risk we have. I would not pretend to say that using an optimizer is going to mean you can throw everything else out. Given the constraints and the inputs that you give it, it'll give you a very good solution, but you then need to do some further stress testing to see if what you have is something that you really want. What I always try to do is get an intuitive grasp of the solution. I try to figure out why the optimizer did what it did. If I can start to think like the optimizer, then I can be much smarter than everybody else that's managing money.

Now we'll try and get into a more specific example. I was talking in so many vague generalities, but now we'll actually see what's going on here. The BondEdge total return optimizer has three really distinguishing characteristics—three things you can really control to get a different type of solution. The first is portfolio level risk constraints. What do I want the duration of the portfolio to be? Where do I want my sectors to be? The second is more practical constraints. If you're actually asking BondEdge to go in there and pick certain securities for you, you may not want it to buy odd lots. You may not want it to buy a lot of any one security. So we have some practical constraints. We need to take into account transaction costs. An optimizer will churn your portfolio forever if you don't calm it down a little bit and realize that there are costs associated with this. You may not want to go through a lot of work trading your portfolio to get another two basis points in anticipated

return if this doesn't seem like it's effective. We have some ways you can practically limit what the optimizer is going to do for you. Finally, it will allow you to use either a single scenario, which kind of gets into turning it into a yield optimizer, or probability-weighted market scenarios. We'll look at each of those in a little more detail.

We utilize a list of portfolio risk constraints that we help people concentrate on. We allow you to specify sector weighting by Treasury, by agency, by all the different corporate sectors, by mortgage-backed securities, and by CMOs. You can really fine-tune those assumptions or those constraints. We also let you limit your credit exposure. You may not want to have anything below investment grade in your portfolio. You may want to have an overall quality of A or better. You can specify that you don't want to have any more than *x*% of the portfolio in anything less than AA and really fine-tune those assumptions. We give you a lot of flexibility with duration ranges. You can put the duration between four and five. You can say you want each sector to contribute *x* amount of duration. You can define the maturity ranges. You can make sure that you are spread out evenly across the yield curves so you're not subjecting yourself to any undue yield curve risk. We let you limit your issuer exposure and your issue exposure. For example you can decide to not have more than 10% of your portfolio in any one issuer, or any more than 2% of your portfolio in any one issue.

What we wanted to do is to take a specific example and give you an idea of what an optimizer would do with it. In the portfolio average constraints I set up, I limited my effective duration to between 3.5 and 4.5 years. I gave myself a fairly wide range, but that's not going to be as important. The convexity constraint is –0.05, which is virtually nonnegative, to just a little bit positive. I was really trying to keep about a zero convexity on the portfolio. Then, in my set distributions, I just did something very simplistic. I needed to have at least 25% of my portfolio in Treasuries, and I can have up to 75%. I said I wanted a minimum of 20% and up to 50% in agencies. I have a minimum of 10% and a maximum of 100% in my pass-through.

There are a lot of other categories I could get into, and I just didn't really bother in this example. I just wanted to do something quick so you can see what's going on. That's really how you would set up your constraints.

There's an issuer issue button, if I wanted to constrain that, which I didn't, but I should have.

I used only one scenario when I ran this. I did an unchanged scenario, and I gave BondEdge a universe of mortgage-backed securities, of agency bonds, of Treasuries,

and of strips, and it found four bonds for me. I gave it thousands, and all it wanted was four, which kind of makes me think that maybe I should think about putting some of these practical constraints in that say I'm not going to put 100% of my portfolio in any one bond. It put 50% of the portfolio in an agency bond, 25% of it in an 8.5% Ginnie, and then in the remainder of the portfolio was put in Treasuries. They're both fairly long Treasuries, so, the overall duration of the portfolio is 4.5 years, which was really the outer range that I had set up, and BondEdge isn't too stupid. Since interest rates aren't moving around, it's going to go out long on the yield curve to pick up as much yield as it can. It gave me a positive convexity of 0.04. If I had been rebalancing a portfolio, this would show me what I had to sell and what I had to buy, but I just started fresh with cash.

We can also take a look at the distributions. It barbells the portfolio with 75% of the portfolio is between two and three years duration, and 25% of the portfolio has a duration that is longer than ten years. I tried to barbell to help it generate some yield for me on the return. The sectors it chose were 25% in Treasuries, 50% in agencies, and 25% in pass-throughs. Those are really nice round numbers which say BondEdge is bumping up against some constraints. I didn't really give any overall quality constraints because all I had was AAA stuff in there. I had agencies and mortgage-backed securities, and Treasuries so, I didn't really worry about my constraints on the quality. Maturity years are also slightly barbelled, and that's about all that's interesting.

It did take very high coupon stuff—everything higher than 8%.

We can compare what the optimizer picked versus the constraints that we've put in to see where I start bumping into my constraints. If I had allowed it to go longer in duration, and allowed it to buy more of the agency, it would have because it bought as much as it could. It didn't buy as much as it could of the pass-through. It bought the minimum of Treasuries. Whatever was left after it bought the agencies and the Treasuries was what it put into the Ginnie Maes.

If I was going to go back and analyze this, I might wonder if it is realistic to have the optimizer pick the high coupons it did. I might want to think about that a little bit. Second, did I mean to constrain this portfolio as much as I did? This 50% agency constraint I put in is really limiting the return I can get. Is that something that I want to put in there or did it have some unintended consequences? Obviously this was a very simple example, but it kind of gives you an idea of what an optimizer can do for you.

The next level of constraints that we can look at is practical constraints, and for these I'm talking about minimum lot size. I don't want it to go out and buy 50¢ of

any one security. I want it to actually be in some reasonable lot size. Then maximize the participation value that I'm going to have of any one security. It's not going to go out there and buy one bond to help me meet my goals. I might want to put some transaction costs in there, so it doesn't churn my portfolio unreasonably. I may want to restrict what bonds it can sell. The example that always jumps into my mind with this is how can you go out there and talk about selling a private placement? An optimizer won't know that it can't sell it unless you tell it not to. You might also want to restrict your sales due to the fact that you have some capital gains that you don't want to take for tax reasons. So you can see what your unrealized gains are and decide not to sell these bonds.

The final thing we allow you to do is to put a yield advantage on your current holdings. This gets back to the idea I was talking about. You don't want to go out and sell all these bonds for another two basis points of return. Let me give a yield advantage to the bonds I have in the portfolio. Maybe I'll give them a ten-basis-point yield advantage. The optimizer won't try and sell them unless they're more than ten basis points below the bond it's going to replace them with.

What I wanted to do is take in a small portfolio of corporate bonds and a universe of corporate bonds and impose these constraints one at a time so we can see how much these things are costing us. There's nothing really special about these bonds.

I made a $5 million maximum participation, which is the first practical constraint I applied.

The second constraint I applied was a $100,000 minimum lot size. The third one was a bigger minimum lot size to see how much that's it is costing me. Then I restricted the sales of two of the bonds in my portfolio just to see what that cost would be. I gave a 20-basis-point yield advantage to my current holdings to see what that would do. The final thing I did was add a transaction cost. Then I ran my optimizer a whole bunch of times, under just one scenario running. I did this a while ago when interest rates were falling.

Of the possible bonds, let's see what the percentages of the portfolio are. My first constraint is my maximum of $5 million of participating value. I applied only the one constraint, and I ran the optimizer. It picked out certain bonds, and it assigned a certain portfolio weighting to them. The total return I got from that was 11.48%. That's what BondEdge says. If you buy those bonds in that percentage and this economic scenario happens, then you'll get an 11.48% return.

I applied my second constraint, which was my minimum lot size, and with that constraint, I had a total return of 11.46%. That minimum lot size cost me two basis

points of return. I applied my other minimum lot size constraint. I made it an even a bigger minimum, and I lost another two basis points of return. This is how I can really get an idea of what all these different constraints cost me. I also restricted the sales of two bonds, so that I ended up with a Mallory bond and a Sears bond in my portfolio that didn't exist before. That cost me eight basis points of return. That was a more expensive thing to do. I also gave myself a 20-basis-point yield advantage for those bonds. I ended up with a higher return. You kind of have to back out what that is, but it does give you a different set of bonds.

The last thing I did was impose some transaction cost. You can go through this kind of analysis to see what your different constraints are costing you. You don't have to use a huge universe of bonds to do it. You can do something fairly compact and concise.

The next thing I want to do is talk about the impact of uncertainty. Everything I've done to this point has been done using one economic scenario. I want to really show you what happens when you let yourself have several economic scenarios. What's the difference between using one yield curve and a variety of market-weighted yield curves? Unless you're very, very good at assigning your probability weights, you're going to see a difference, and the difference is the convexity or optionality in the portfolio.

When you let yield curves change, Treasury rates move around, and the options become more or less valuable on the bonds, and any optimizer is going to look to get more convexity. It's going to say there are certain bonds that do respond better when interest rates move, and that's what the positive (or less negative) convexity is all about. If you have one economic scenario, it's going to be yield positive. It's going to go out there and say I don't care what happens to these options because there's no volatility out here. Nothing's ever going to change. I'm going to go out and buy all the yield I can get. When you move things around, then the bonds that have higher convexity become more valuable. Adding the uncertainty changes what you're going to get for a solution, and it certainly impacts the option-embedded bonds.

To try to illustrate this, I didn't even have to pick up any option-embedded bonds. I just wanted to look at a very simple scenario. My first scenario says I'm going to have a no-change interest rate scenario. I'm 100% sure that's going to happen. My second scenario says rates are either going to go up 100 basis points, stay the same, or go down. I've centered my probabilities around the fact that they're not going to change, but I know that there is some potential volatility out there. I'm 50% sure they're not going to change, but there is a 25% probability they will go up or a 25% probability they will go down. When we have more than one scenario, what is the

optimizer going to try and do? What is it trying to maximize? It's trying to maximize your expected return. For each bond, and for each economic scenario, it's going to calculate what the total return is over your investment horizon, and then it's going to probability weight them. What the optimizer's going to use is just the expected return that you're going to get. For actuaries that's a pretty easy concept, but it isn't for some of my clients. I'm not trying to minimize the volatility of returns; I want to maximize my expected return.

In my example, I gave myself a very simple universe. I just picked four on-the-run Treasuries, and I said my duration is going to be five. I didn't really let the optimizer pick. Then I used two scenarios. For the first scenario, it picked mostly the five-year, on-the-run Treasuries and some of the ten-year, on-the-run Treasuries to get the duration right. It gave me a yield of 6.26% with a convexity of 0.16, and we'll talk about that nonparallel duration later. The second scenario doesn't have option-embedded bonds either. There's nothing that's going to change in value, but it's looking for more convexity. When interest rates change, more positively convex bonds do better. It's going after the 2-year and the 30-year because that combination gives it 0.41, and I gave up 10 basis points in yield, but my expected return is going to be higher than if I had just gone with the 5 and the 10, if I'd done a more bullet strategy rather than a barbell strategy.

I picked a nonparallel duration number, which is something I think unique to Capital Management Sciences. What we try to do is assess, in one number, your sensitivity to changes in the shape of the yield curve. To calculate a nonparallel duration number, I leave the long end of the yield curve where it is, and I move the short end up and down 100 basis points, and I calculate my price sensitivity to that. Under the first scenario, I had 1.26 as my nonparallel duration. I always like to look at this as a relative number. Under the second scenario, I really decreased that nonparallel duration number down to 0.82 which means that I'm really positioned for a flattened yield curve, and if I don't think that's going to happen or if I didn't want to position myself that way, that might be something else I want to look at. I hadn't tried to do anything in terms of optimizing what I thought the right nonparallel duration number was, but this fell out of the analysis, and it's something I now I need to think about. Is this what I really intended to have happen? As I said, the optimizer is cold and cruel—it doesn't have any subjective thoughts. It's only going to do the analysis based on what you tell it, and that's what it's coming up with. It really does change my exposure to a change in the shape of the yield curve.

To conclude, total return optimization can be used in three different ways. It can help you generate some ideas and work on your portfolio strategy. It can also help to identify opportunities and help you invest funds. You really need to analyze and

refine these "optimal" portfolios because when they come out of the optimizer they're not necessarily optimal. You need to do a little bit more work and think about what they've done to you.

Mr. Lawrence A. Berger: The model I'm going to talk about has actually been developed in a property and casualty context, but with some work, it can be adapted to the life insurance side. It actually came out of a Ph.D. thesis written at the University of Pennsylvania in 1978. If anyone's interested in any references, there have been a number of articles written on this model.

Before I get into the model, I'd like to mention another approach to this general question of asset/liability management and portfolio optimization, and that is the techniques of stochastic optimization. I don't know if anyone is familiar with the techniques that have been developed in operations research, but they're quite useful given that you already have a well-developed simulation model which, of course, many life insurance companies do have. A leading expert in the area is Professor John Mulvey of Princeton University. He has worked with a number of companies such as Towers Perrin and Falcon Asset Management, besides many others in other industries. He has quite a history in this regard.

The stochastic optimization technique does require significant investment in computer technology, along with a well-developed simulation model. This is one thing to keep in mind. On the other hand, the approach I'm going to talk about also relies on having a well-developed simulation model of your business. From that point on you can implement this model in an Excel spreadsheet. What it does is essentially extract key information from the simulation model, and then you can put it into an Excel spreadsheet and vary certain parameters on a much more global basis. I'm not pretending that you can model fixed income in detail with a model like this, but rather you can look at line-of-business classifications on the liability side and asset categories on the asset side. Then you can see what the total corporate risk and return profile is for a given configuration and a given level of capitalization.

Asset/liability management is a pretty hot topic in the property and casualty industry these days, but it goes by the name of dynamic financial analysis (DFA). DFA provides the user with a way to manage the total corporate risk and return. Now, one good reason to go through an exercise like this is it gets you to focus on some basics. Many companies in the property and casualty industry have not looked into, for instance, what sort of target return you have for the corporation as a whole. That's the first step you need. Given that you have these targets, what level of risk are you willing to live with?

One benefit is it gets you to start focusing on the basics that most corporations think about.

Once those issues are settled, you can ask a set of questions that are really interrelated, such as how much capital do you require? What level of capitalization do you need to have? How are you going to invest your assets? What are you going to do on the line-of-business side? Does it make sense to alter your line-of-business mix? What this points to is the fact that when you look at an insurance company, you're really talking about a double portfolio problem. On the one hand, you have the standard asset portfolio problem, but you can also view your line-of-business mix as another portfolio. These can really be viewed as two subportfolios that are aggregated to form a global portfolio, and that's exactly what we're doing here. We're using standard mean variance technology, except that I'm adding in skewness. I'm adding in a third moment to the standard technology, but I'm looking at combining these two portfolios into one global portfolio, and then tracing out an efficient frontier where the measure of risk is not standard deviation, but is probability of ruin or probability of surplus decline.

So, again, we have the questions of what asset mix, what line of business mix, and related questions on the line-of-business side. A radical way to alter your line-of-business mix, of course, is to think about an acquisition. Another way to alter your profile in your line-of-business side is to alter the reinsurance program.

When you have a simulation model, and you want to check what happens when you change these various dimensions, you have to run that simulation model over and over again as you change the settings in the various dimensions. The key idea here is that with this type of approach all you need to do is essentially run the model once and extract key information from it. From there you can go into an Excel environment and plug in different values for your asset mix, for your line-of-business mix, and for your leverage. Now, that's not true when it comes to reinsurance. Because of the subtleties of reinsurance arrangements, you have to go rerun a model if you do change the reinsurance program; otherwise it does significantly reduce the amount of time you spend running simulation models. Given that you can tell what your risk-and-return profile is for virtually any configuration, you can ask Excel to search within constraints over the universe of possible configurations to come up with, in this case, for any given level of risk what the best return is and what configuration provides you with that best return. I will discuss an example where we're going to choose a 0.25% probability of surplus decline of a certain magnitude and then determine what level of return you can earn as a corporation and what the configuration would be. Then you can go to 0.5% probability of surplus decline and see that you will get the expected increase in return.

The most important possible limitation when it comes to the life side is that you do need to be able to separate underwriting from investment results. That's fairly straightforward when it comes to the property and casualty side. It may be more difficult on the life side. I actually developed this model when I was at M & R, and we talked, albeit not very seriously because we didn't have a client paying for it, about the possibility of applying it on the life side. It's not clear, but I do think that it would be possible, with some work, to modify this. I will be explaining an example worked out on the property and casualty side, where we talk about underwriting results primarily in terms of loss ratios, which is simply the ratio of losses to premiums.

What we need to do is start with some distributions for each line of business to try to get an assessment of the risk and return associated with each line. We're going to do it, of course, net of reinsurance, and we want to try to reflect all the risks that are associated with writing these lines, including any loss development experience and reinsurance collectibility. For example, let's get a little concrete and talk about a company writing three lines of business: medical malpractice, workers' compensation, and general liability. I don't mean to make it look easy. There's a lot of actuarial work that would have to go on, and it has gone on. These models have been applied for a number of clients. It takes a lot of work to come up with an expected loss ratio, a standard deviation, and a skewness. Coming up with that third moment is quite important on the property and casualty side, as you would imagine, given that it's subject to catastrophes of various sorts. In my relatively short tenure on the life side, I realize that you have plenty of downside risk here, too. For instance, with the guarantees that are built into a lot of products, there's a lot of downside risk. I think this issue of downside risk is quite significant on the life insurance side, too, and we are including that.

That is a significant limitation of the typical portfolio optimization models that only look at mean and variance.

We do that work to get the loss-ratio distributions, but that's not good enough. We can't just look at these lines of business in isolation. We need to understand how they interact because we want to try to get up to the corporate level. What we need to do is combine these loss-ratio distributions and come up with an aggregate distribution. How are we going to get at the linkages between the lines of business? The best way, in my view, is to develop a simulation model that has those linkages. What are some of possible sources of interrelationships between the lines of business? Of course, we know about the systematic influences that the economy brings to the line-of-business result. You're going to need to include changes in interest rates, inflation, and other variables in the simulation model. There can be

dynamics in insurance markets that also create interactions between these lines of business. You need to model these and then extract the key information.

Portfolio theory tells you how to arrange those lines of business. Given any set of components that form a portfolio, we know that the covariances are going to determine the linkages that enable you to get to the aggregate level to determine what the variance of a portfolio is. For instance, I'm sure you're familiar with the equation from portfolio theory that tells you that the variance of a portfolio is essentially the sum of the variances plus the sum of covariances between the components of the portfolio, weighted by the product of the percentages. Similarly, if you want the skewness of a portfolio, there's an equation that looks exactly like it, except it's a triple sum, and you have a coskewness and three products of the percentages. That enables you to get to what the aggregate portfolio risk is given what these individual covariances are.

Once you know the covariances, you can change the allocations and the percentages in the various lines. It is just a question of plugging in numbers, and that's really what this model does. We calculate covariances and coskewnesses, and once we know that, all we have to do is plug values into an Excel spreadsheet, and it will give us the overall standard deviation and skewnesses of any given portfolio. But we do need a simulation model in order to produce the estimates of those covariances and coskewnesses. That was just the loss ratios. We want to get to the total underwriting result. We'd like to model all underwriting risk, so we have to include premiums and expenses. Usually, you lose money on the underwriting side, and hopefully you make money on the investment side. Underwriting profits that are going to be, in property and casualty lingo, 100 minus the combined ratio. The combined ratio is the ratio of losses and expenses to premium.

We have a risk profile, on the underwriting side, of profits, standard deviations, and skewness. Then we need a similar profile on the asset side. That's what I have next. I have five asset classes, and I don't pretend to be able to model U.S. fixed income in detail.

I went through lots of exercises regarding parameter uncertainty and how robust the model is to uncertainty about things like correlations. You find that because of the high correlations within the fixed-income categories, a model like this is quite limited in discriminating between fixed-income classes, but I do believe it's quite good at thinking about your equities, your allocation between equities and fixed income, and possibly international investments.

We have the components of the asset portfolio, so now we need to put the whole package together. How do we do this? For that, we again need a simulation model. What we need to think about is not only the linkages between the different assets and liabilities, but also the linkages between the assets and the liabilities. If you're going to go to the total portfolio level, this is what you really need to do in order to represent the risk of the total corporation. Again, you need a simulation model in order to study what those linkages are, but I think that's just good information to know if you're trying to understand what the risk profile of your total corporation is. How does the whole package fit together? That really is what asset/liability management is all about. A model like this really does illustrate the key relationships that exist. If you change something on the line-of-business side, it affects what you should do on the asset side.

With that model, you'll simulate a large number of times, possibly 10,000 or whatever number of times you need for convergence. Each economic scenario has an associated set of underwriting results and asset results. We run the simulation model, and from there we compute the means, the covariances, and the coskewnesses. We extract this key information for all lines of business, for all assets, and between the lines of business and the assets. Once we have these data, we can now calculate the mean, the standard deviation, and the skewness for any line-of-business mix, any asset mix, and for any leverage. We have the first three moments. What do we do from here? From here we can now fit a distribution to it. A similar technique was applied in an article called "Application of Risk Theory to Interpretation of Cash-Flow Testing Results," that appeared in The *North American Actuarial Journal*, Vol. 1, No. 2, by Edward L. Robbins, Samuel H. Cox, and Richard D. Phillips that used this exact technology, taking a simulation model, and extracting it. During the first iteration of the referee process they used the exact method I used, which is the method of moments on the first three moments. The referees were academics, and they got on them to get more sophisticated. The authors mention, in the article, that the first time through they did just look at the first three moments and used that observation to fit a probability distribution. That probability distribution is an aggregate, and not just a loss distribution. I'm sure you're familiar with the notion of aggregate loss distribution from risk theory studies, but this is an aggregate surplus distribution. It does combine your loss results and your asset results into one probability distribution.

For example, say that company has $1 billion dollars' worth of surplus and writes those three lines of business. Suppose 25% of the premium is written in medical malpractice, 35% is in workers' compensation, and 40% is in general liability. Suppose the company is conservatively invested with 90% in U.S. fixed income and 10% in equities. We also need some information on the leverage of the company. We can assume the company has a premium-to-surplus ratio of 0.75. We also need

the reserve-to-premium ratios for each line of business because different lines of business are going to pull in different amounts of investable funds. These reserve-to-premium ratios give you the linkages between the assets and liabilities. If you change your line-of-business mix, you're going to have more or less in the way of investable funds, and the whole portfolio will change. That links together the assets and liabilities, and given a premium-to-surplus ratio and these leverage ratios for each line of business, you get the reserve-to-surplus ratio.

We run the simulation model. We extract the covariances and coskewnesses and come up with the expected risk profile. The return on surplus that is expected in this example is 14.37. We have a standard deviation of 20.24 and a skewness of 1.74, so let me point out that you cannot plug in any other values here. You're in an Excel environment where you can just change those line of business percentages, change the leverage, change the asset percentages, and get this off of the spreadsheet. We can take the first three moments and fit a distribution.

You can fit any distribution, but I fitted three distributions that allow you to have some nonzero skewness. For comparison, I also determined what happens by fitting it to the normal distribution. That showed me how important it is to use a skewed distribution in this context. In this case, I was just looking at the aggregate distribution and chose the point where I had a 50% decline in surplus. When I assumed normality, I got virtually no chance of a 50% decline in surplus. The other distributions come very close to each other. There's not a lot of difference between these distributions, if you know the first three moments. I'm not trying to minimize the difficulty in reaching an assessment like that, but given that you have information like that, there's not much difference between the estimates. All these estimates say that in this particular example there's about a 0.5% chance, or a 1-in-200 possibility that there will be a 50% decline in surplus. So the company would be considered to be in good shape.

What we can do now is ask Excel to search. Excel does have an excellent solver in it, and you can beef it up. There's a system available from Front Line Systems that can beef up that solver to get you a pretty good optimization routine. You can throw all sorts of constraints on here. For instance, let's just vary the asset mix. After this I'll vary the line-of-business mix also. We have 10% in U.S. equities and 90% in U.S. fixed. I'm going to also set the risk level equal to 0.25% probability. The current position is that there's a 0.5% probability of a 50% decline in surplus. I'm going to reduce the level of risk. You could again throw on other constraints because you're in a spreadsheet environment. For instance, suppose you want your risk-based capital ratio to be bigger than or equal to two. You can throw that in as a constraint, too.

It's quite a flexible environment. What's the result? Excel searches and it finds an asset portfolio that will get you a 17% return on surplus. It was 14.37% with a 0.5% probability of a 50% decline in surplus. We come to 17.05%, but the probability of surplus decline has gone down to 0.25%, and what's the associated asset mix? It finds that you should be more diversified. It doesn't take a genius to figure that, although you might not look so smart, especially if you invest in international equities. The good thing is that you can incorporate downside risk. You can throw in scenarios. I was throwing in adverse scenarios just to pick up the fact that there can be big downside exposures, and the coskewness picks that up. That goes into these parameters that estimate what the skewness is of the portfolio. So this is the benefit of including downside risk in an optimization model like this.

The next thing I've done is set the constraint equal to 0.5%. Now we're asking, what if we live with a higher level of risk? Then the model says that you can get up to 19.5%. It was 17% at the lower level of risk of a 0.25% probability of surplus decline.

Finally, you can also vary the line-of-business mix along with the asset mix. The interesting thing here is that you'll see that when you change the line-of-business mix, what you should do on the asset side changes, too. This is a pretty clear demonstration of this basic fact of asset/liability management. It's considerably more difficult to change your line-of-business mix, so I've set some relatively tight constraints that Excel is allowed to search over. However, an acquisition could significantly change the composition of the line-of-business portfolio.

The general rule is, as you loosen constraints, the further out your efficient frontier goes. In our example, we get close to 20% at the 0.25% probability of a ruin level.

The model says to invest more in medical malpractice, invest less in workers' compensation, keep general liability the same, and change the asset mix accordingly. If we go to the 0.5% risk level, and you vary the lines of business and the assets, we get up to 21.5%.

From the Floor: This is ground zero for me in terms of understanding what's going on, but I have a general question with regard to the optimization black box. If you run the optimization, regardless of what you're trying to determine, will multiple runs yield the same portfolio construction from a universe?

Ms. Ehrlich: Yes. There's nothing random about the solution. If you start with the same set of bonds, the same constraints, and the same scenarios, it will give you the same solution. It's just a linear programming routine that's behind it.

From the Floor: I have a question for each of you. Cathy, I didn't understand this last business about giving yourself 20 basis points on your current portfolio.

Ms. Ehrlich: What I'm trying to do by giving myself a yield advantage is to somehow trick the optimizer into thinking that the bonds that I have in the portfolio are actually going to yield a little bit more than would ordinarily happen. As I said, BondEdge is going to be very cold, cruel, and heartless. If some bond is going to give me one more basis point on return, it's going to tell me to sell what I have and buy that bond. There are not too many people out there in the world that are going to go through that transaction just to get one basis point of return. By putting that 20 basis points in, I'm going to pretend that all the bonds in my portfolio are just 20 basis points better than they really are. So the standard to sell those bonds is higher, if that makes sense.

From the Floor: So, it's partly to avoid transaction costs and churning.

Ms. Ehrlich: Yes.

From the Floor: When you're doing your projections and simulations, are you simulating for just one year?

Mr. Berger: For one time period. On the life side, I think it might be more important to look at a five-year period, but it does have to be one period. That is a limitation, but it doesn't have to be just one year.

From the Floor: You had ranges for your lines of business as well as for your asset allocation.

Mr. Berger: Right.

From the Floor: Is it trying to tell you what kind of business you should write?

Mr. Berger: That's right.

From the Floor: What are you trying to optimize? What's the criterion to pick different percentages for the lines of business? I assume you're trying to maximize your yield for the assets.

Mr. Berger: This is a total corporate model. What we've developed through all these stages is a total return on surplus that includes your returns on your assets and liabilities. It doesn't matter whether you're tinkering with the asset side or the liability side. You're doing the same thing. You're asking what is the highest return

that can be achieved given a certain level of risk? What's the configuration that gets me there? I want to point out that you could have varied the leverage, too. Excel can search over the surplus. I was really surprised to see that it comes up with numbers that make sense. The numbers do not vary too much, and as you increase your risk level, the leverage goes up, as you would expect, but you don't get corner solutions which is something I thought you might get.

From the Floor: Well, I assume you can run this twice—once where you fix the business allocation. It seems that sometimes you don't have too much choice about what business you have.

Mr. Berger: Right. In the Excel environment, you just plug in numbers for what your constraints are, and then it runs very quickly.

From the Floor: I want to ask Cathy if she foresees the usage of more sophisticated techniques of programming, like, for example, genetic algorithms for optimization purposes.

Ms. Ehrlich: That's a very big question. I've been to a number of conferences, a number of papers have been written, and a lot of people are talking about it. I've actually seen some applications where it appears to work very well for equity markets. I am not a visionary, but what I have seen in the fixed income world hasn't really convinced me that this is the direction we'll go, but I wouldn't bet on anything right now. That's just what I've seen. The reason that it appeals to me is that I think in the fixed-income world, there are very good ways of capturing that risk and of understanding how that market works. As I said earlier, if you do a principal components analysis to try to figure out what makes bonds return the way they do, so much of it is determined by very easily understood factors. I'm not sure how much is gained by all that other stuff I don't really understand, but there are people smarter than me who are working on it.

From the Floor: I have a comment on using the genetic algorithm. There will be a paper that will be published in the *Journal of Actuarial Practice,* Absalom Press, Lincoln, NE. The subject is an asset/liability model, so it generates the surplus cash flow. You have your liability, such as a single-premium deferred annuity or whatever, and then the model sets asset allocations and will give you certain asset allocations. It will give you your return on surplus and all that, and then it also gives a simple, genetic algorithm so that you can search the asset allocation that will give you the profile that you want. That paper is just a real simple, genetic algorithm. There needs to be more work done to get the best genetic algorithm. So, look for that paper.