RECORD, Volume 24, No. 3
New York Annual Meeting
October 18–21, 1998

Session 8PD
Equity-Indexed Products: Managing the Beast

Track: Investment
Key Words: Annuities, Asset/Liability Management, Asset Topics, Equity-Indexed Products, Models

Moderator: ANSON J. (JAY) GLACY JR.
Panelists: ANDREW S. CHOW†
ANSON J. (JAY) GLACY JR.
JEFFREY J. LOBO‡
STEPHEN J. STONE
Recorder: ANSON J. (JAY) GLACY JR.

Summary: This panel discussion focuses on the problems encountered in managing a book of derivatives supporting an existing equity-indexed portfolio. The different approaches companies have taken are covered, including: over-the-counter versus exchange-traded contracts, using the "greeks" to measure exposure, dynamic hedging, and managing marketplace volatility.

Mr. Anson J. (Jay) Glacy Jr.: We've assembled a first-rate panel and program to explore the hedging issues associated with equity-indexed products (EIPs). I'm a senior consulting actuary with Ernst & Young (E&Y) based in Hartford, Connecticut. I'm part of E&Y's risk management team, focusing on the capital markets risk exposures of life and health companies. My job today is to describe dynamic hedging, define the "greeks," and talk about liability modeling issues.

Our second speaker will be Steve Stone, director of global derivatives for Allstate Life. Steve will take us beyond the greeks, showing us their shortcomings and where they lie.

---

†Mr. Andrew S. Chow, not a member of the sponsoring organizations, is Vice President, CONSECO Capital Management in Carmel, IN.

‡Mr. Jeffrey J. Lobo, not a member of the sponsoring organizations, is Vice President, Keyport Life Insurance Company in Boston, MA.

Following Steve will be Jeff Lobo, vice president of risk management for Keyport Life, who directs the risk management activities on the investment side. Building on Steve's ideas, Jeff will illustrate how one might optimize dynamic hedging strategies for EIPs.

Finally, Andy Chow, vice president of CONSECO Capital Management (CCM), is the portfolio manager with responsibility for all derivatives activity at CCM. Andy will tie all the pieces together with a case study.

As I understand it, all of the top indexed product writers use some form of dynamic hedging to manage their market exposure. Dynamic hedging entails the establishment of positions in certain instruments whose market movements parallel those of the liabilities to be hedged. For indexed products, the following instruments typically are used: (1) Standard & Poor's (S&P) 500 futures contracts, (2) S&P 500 index options, and (3) some Treasury futures. These would be layered on top of your bedrock bond position.

What are these mysterious greeks? If you know duration and convexity, then you know the greeks. They are essentially measures of sensitivity to change in some underlying base. Delta is the primary greek. You've probably heard the term "delta hedging" before. It's the first-order change in an EIP's value with respect to the underlying index, in this case the S&P 500.

Gamma, the counterpart to convexity on the interest rate risk side, is the second-order measure, the rate of change in delta with respect to the underlying base. Next we have vega, which is an honorary greek because it is not a greek letter. Vega measures the rate of change in an EIP with respect to changes in market-implied volatility. Theta is the rate of change in value with respect to time. Finally, rho, the true counterpart to our familiar duration, is the rate of change in value with respect to interest rates. For long-dated contracts, like EIPs, vega and theta usually exceed in importance what the textbooks say. And recent marketplace events teach us never to ignore vega. My partners on the panel will talk in more detail about vega.

What are the key features of a facility for calculating greeks on the liability side of the balance sheet? First is stochastic mechanics. It's critical to have a two-factor model that suitably traverses the event space and covers both movements in interest rates and the equity market. Second, I think that we're starting to understand that policyholder behavior for indexed products will be different from their behavior for conventional products. Policyholders may have what we call a "naïve" (buy high, sell low) and a "savvy" (buy low, sell high) mode, two different philosophies that might explain policyholder behavior. The modeled correlation between the movements in the equity markets and the debt markets will be critical to predicting such activity.

Of equal importance in modeling and computing greeks is the formalization of a renewal rate crediting strategy in response to market dynamics. Essentially, I think of this as an artificial intelligence algorithm for predicting how renewal participation rates can be modeled. I have two examples of what we call "adaptive" mechanisms that show how the future participation rate in the contract might change:

- $PR_t = (BS_0 * PR_0) / BS_t$
- $PR_t = (BS_0 * PR_0 + ER_t - ER_0) / BS_t$

The first example very simply changes the participation rate in some future year in response to the proportionate change in the corresponding Black-Scholes option price. A more complicated adaptive participation rate formula also adjusts future participation rates for changes in the asset-earned rate that the insurance company is realizing. Finally, on the liability side, it is critical to recognize multiple index terms. Don't stop your calculations after one index term. The key is to model when the policy ultimately will turn into cash and how much.

Mr. Stephen J. Stone: Implied volatility is a way of calculating volatility by backing out every single input into an option-pricing formula, except the volatility. The advantage of looking at implied volatility is that it's a standardized measure that lets you compare across strikes and time.

One critical attribute of implied volatility that's worth mentioning is how volatile it is. The Chicago Board Options Exchange (CBOE) publishes an index on a real-time basis called the VIX Index, an index of one-month, at-the-money implied volatility. On October 5, that index went over 60% for the first time since October 1987. Less than three months before that, the index was trading below 20%. The comparison gives you some feel for how volatile implied volatility and option prices can be.

Note that implied volatilities aren't the same for all options, and it's important to understand some structural aspects to be able to price and hedge any options that are embedded in equity products appropriately. The first attribute is the volatility skew. I took prices from the CBOE for March 18, 1999, and maturity options as of July 31 and calculated their implied volatility. For 900 strike options, the implied volatility was 29%, and for the 1,400 strike options it was around 18%. So, there's a significant difference in the implied volatility between low-strike options and high-strike options, with low-strike options typically being puts and high-strike options typically being calls.

There are a number of reasons for this supply/demand phenomenon. There's a demand for out-of-the-money puts, so their prices are higher. People use those for hedging purposes. Another important structural aspect of implied volatility is term structure. As you move to longer expires, implied volatilities increase. At the end of July, implied volatility on near-term options was below 20%. But if you looked out at the December 2000 Long-term Equity Anticipation Securities (LEAPS), it was up around 23–24%. I'll speak a little bit more about the term structure and the skew later, but first let me mention that there isn't any single number that you can point to and say, "This is objectively what the delta or gamma of this option is."

Let's look at a case study. On September 4, XYZ Insurance Co. wants to hedge its equity-indexed annuity (EIA), which is a 1-year 90% participation product with a 12% cap. I've set the numbers up to come out exactly to 100,000 S&P units because the S&P was at 973.89 on September 4. The 1-year, at-the-money volatility is 32.2%, and the 1-year, out-of-the-money volatility is 28.3%.

I calculated the greek in two ways, using methodologies that sophisticated market participants might use. The volatilities are drastically different. If I were using implied volatility to calculate my delta for the call spread, I would purchase 69 S&P futures contracts and say I'm delta-hedged. If I were using a level 15% volatility, which is the long-run level of volatility that tends to be exhibited by the S&P, I'd be buying 125 futures contracts. I'm simply dividing those numbers by 250, which is the number of S&P units in a futures contract. If the market were to rally or to sell off, the profit and loss (P&L) of the dynamic hedger using these two methodologies would be substantially different.

Using implied volatilities, the gamma and vega are positive, but when using the level 15% volatility both are negative. If I wanted to hedge and I were using implied volatility, I'd be selling options. Someone using the 15% volatility would be buying options to hedge. That's a drastically different hedge. Furthermore, if you weren't using options to hedge, which is delta-hedging, the gamma hedger using implied volatilities would be selling S&P contracts as the market rallied and buying them as the market went down, but the person using a 15% level volatility would be doing the opposite. The major point I'd like to make is neither of these is the right way to do it, and it really depends on somebody's objective.

Here are two examples: If your objective is to minimize tracking error on a daily basis, then you're probably not going to do any better than using implied volatility to calculate your greeks. However, if you weren't concerned about tracking error and wanted to put the cheapest possible hedge on, you wouldn't be putting any hedges on at all. That's not to say that writing out-of-the-money puts naked is a good idea, but the basic point is that, depending on your objectives, you're going to come to dramatically different results for constructing your hedges. When you talk about the greeks there is no one objective number, which is the delta of the option. The greeks are ambiguous.

I put together a case to illustrate the problem of just using the greeks to hedge. I used implied volatility to construct the hedges. I won't go through all the details, but it's similar to the previous case: We have a 90% participation rate, a 12% cap, and $1 billion of EIAs. We're using S&P 500 futures contracts to hedge delta and Eurodollar futures to hedge rho. Eurodollar futures is one way to hedge rho, as are Treasury futures. And, because we are using implied volatilities, we are a seller of options to hedge gamma and vega. We sold 1,500 of the 1,125 calls of September 1998 and 400 of the 1,125 calls of June 1999.

What happens during the next two weeks? We adjusted our futures contracts on a daily basis to make sure we were always delta- and rho-hedged and, because we did that, essentially, at the end of the period, we broke even on delta, gamma, and rho. However, we experienced a $1.5 million vega loss on the liabilities and a $700,000 vega loss on the hedges, which is somewhat curious. We put these hedges on to hedge vega and they lost money along with the liabilities. Something seems to have gone wrong. The total loss we experienced was $2.2 million entirely due to vega. Whether that's a large number or not for your firm depends, but I wouldn't want to report a $2.2 million loss.

What went wrong? First, for this two-week period, implied volatilities increased. The VIX went up about 5 points, from 26.27 to 31.08, which explains why we lost money on vega on our hedges. But, we were short vega on our liabilities, so we should have made money if volatility went up. What happened? The volatility term structure also inverted, and, more important, the volatility skew widened. The differential between out-of-the-money puts and out-of-the-money calls widened dramatically. As a result of these shifts, the liabilities were long vega, but we were short the skew because we were short low-strike options and we were long high-strike options. That's being short the skew.

If you decompose your vega gains and losses into different components, which is a sophisticated analysis, the results are fairly intuitive. We actually made $800,000 from the nonskew movements in vega. If the volatility surface had moved parallel, we would have made money, but it didn't and we lost $2.3 million because of the widening of the volatility skew—or a $1.5 million overall vega loss. The hedges lost $700,000, bringing our overall loss to $2.2 million, about the same as we lost through the volatility skew widening.

The vega hedges performed their job very well. We hedged vega, and, following the general movement of implied volatilities, we would have made money. We lost money on the hedges and made money on the liabilities. That all worked just fine. Our vega hedge performed well, but we didn't hedge the skew, which widened substantially. The major point here is that vega is not a sufficient statistic for measuring risk because of changes in implied volatilities. There will be substantial swings in vega. And the skew continued widening after August and has remained at near-unprecedented levels.

Let me summarize why the greeks are inadequate. They don't capture all of the risks that might exist. Vega is probably the prime example. The skew and the term structure are very volatile, and implied volatilities themselves are quite volatile. If you don't have those specific risks hedged, you can subject yourself to a great deal of risk unwittingly. That's what the case study focused on. You can't just say this is your greek statistic; it depends on how you're calculating it, and that depends on your objectives.

Let me make a couple of additional points about interactive effects. As implied volatilities change, your deltas and gamma will change. These can have very significant effects on your hedging program. If your intent is to be delta-neutral, and implied volatilities are moving, you may find yourself buying and selling far more futures contracts than you might have anticipated.

Higher order risks are important, particularly with gamma. Gamma is the second derivative, delta being the first derivative, of the price with respect to the underlying factor. Especially if you have a book with call spreads or hedges on longer EIAs, where you're long and short options, the third moment can become very important.

Direction can have a large impact on what happens to your gamma, and that is not captured in the greeks. An extremely important point is the correlations between the risk factors. As the underlying factor changes, implied volatility will change. But you can't assume that, if the market goes down or up, you're free from risk on your other factors. Specifically, I calculated the daily correlation between the CBOE's published VIX index and the changes in the S&P 500 over the last year. That correlation was minus 0.91, meaning that, as the market declines, the volatility of options will increase dramatically. As the market increases, volatility will go down. That can have a profound impact on how you might want to run a hedging program.

Furthermore, there are substantial correlations between the various factors in the implied volatility surface. Over the last year, for example, the term-structure correlation with one-month volatilities has been minus 0.95. The skew correlation with the term structure has been minus 0.49 and the short-end volatilities and the skew have had a correlation of 0.61. That means, as implied volatilities go up, your skew is going to widen. And that can have dramatic impact on your P&L as a dynamic hedger. To ignore these effects means you can be taking far more risk in your portfolio than you may have intended.

Greeks are model- and assumption-dependent. I've demonstrated the assumption dependency but, as you start getting into exotic options, there are different models with the same input assumptions that might produce different results. So it's important to have a solid model, one that's been tested and one you're comfortable with, for constructing hedges. When the option expires, you're going to have cumulative gains or losses on your hedges. If your model was poor, you could certainly have a problem with the hedges being constructed inefficiently. There's path dependency in hedging results, and this does not apply just to exotic options. Short-term calls can have extreme path dependency in their results. A minor point concerning parallel shifts and rho—parallel shifts in interest rates never happen.

Furthermore, in a portfolio context, the greeks lose a lot of their usefulness. If you start adding greeks in computing a gamma statistic for a portfolio with longs and shorts and think that this is how delta's going to change, that can be very misleading. If you have a 2% move in the underlying factors, and we've seen plenty of those in the last couple of months, your projected delta, based on your gamma, can be way off.

My real intent was to demonstrate that the greeks are very important numbers. I don't mean to downplay their importance. If you just needed to see seven numbers about your options, you could look at their value and the six greeks. However, you need to go beyond those seven numbers to see how they were calculated, what they mean, and what your hedging objectives are.

Mr. Jeffrey J. Lobo: Steve has done a good job of showing how difficult it can be to hedge when you're relying on a number of different theoretical measures of risk. It gets more difficult when you introduce other real-world factors such as various time horizons. We've been thinking about some of these problems at Keyport. There probably is no final answer, but I'm going to share some of our thinking about how to take into account some of these complexities. I'll briefly outline the problem, give you a quick example, tell you where we think some of the difficulties lie, and then discuss the benefits that are motivating us to pursue this.

Steve went beyond the normal textbook example of dynamic hedging, which typically contemplates selling a call option and then using futures contracts to hedge the delta, adjusting your portfolio dynamically on a regular interval to replicate the option. If you do that often enough, you'll end up with a very good match between the P&L on the two portfolios. It's also simple in a lot of respects because you're just looking at delta and just using one hedging instrument.

In real life, however, you might be long and short in a portfolio of options or be hedging a product with an annual discrete look-back, which is a very complex product when you put a portfolio together. On top of that, we're contemplating seven greeks. And there are also the cross products of the greeks—you can make up as many as you want, depending on which ones you think are important. In reality, you're trying to match two multidimensional surfaces over time, and it's a very complicated problem.

How do you keep the greeks matched? There are an infinite number of instruments you have to trade against, so in reality you have a lot of flexibility in the matching. How do you limit this problem to be able to say you're doing the best job you can? I present it as an optimization problem because I'm trying to develop a framework for comparing one solution with another and determining the one that is best for your needs in your particular situation.

In the optimization problem, deciding what's an objective and what's a constraint is almost interchangeable. In a typical optimization problem, you set up an objective function to minimize or maximize, say, the expected cost of hedging the option over time. You're not worried about whether there are some outliers. At the same time, you might introduce some constraints on the process so that, say, you never lose more than 5% of the premium in the worst case possible across a number of scenarios.

There are many ways to formulate this. We've looked at expected hedging cost and tried to minimize the amount of cash that we have in our option positions. One thing that people don't normally talk about is that you can make these products a lot more or less profitable simply by keeping money in a bond portfolio earning a higher spread in exchange for credit risk.

You might also want to minimize your capital charges. You might want to make your return on regulatory capital the best that you can. You can think of many return-related measures. We also spend a lot of time with risk-related measures such as asset shortfalls, either during the index term or at the end of the term, or mismatches on the greeks. We typically look at the greeks, something that we believe captures the most important directional risks in the portfolio, and then at the liquidity position and limiting the size of positions.

One of the things we concentrate on is using European call options in our portfolio rather than buying highly customized options. Partly we do that for liquidity. Another thing that we have been pretty successful with is limiting the number of transactions that we enter into, which can substantially affect how profitable these products are. So, you could incorporate all of these things into an optimization strategy by using one of them as the objective and the rest as constraints. You can also think up some new ones. A lot of flexibility is possible when applying this kind of approach.

We've decided to make the problem discrete. We look at a number of discrete market states over time and generate a set of paths that I believe characterizes the market. In your particular situation, you might want to use the equivalent of the "New York Seven" or some other type of judgmental scenario set. We've calibrated a lattice to option prices in the market for reasons of internal consistency, which I'll get to later. Then, at each node in the lattice, we think about what trades we can do. You can define an arbitrarily large number of trades, depending on how much you want to track. What options can we buy? What futures can we buy and sell? You might also have another dimension consisting of the implied volatility of the options or interest rate levels.

The third thing we do is calculate values for all the available trades for the liability at each one of these states. All of this becomes input to one big optimization problem. These are all the choices I have available in my restricted choice set and these are the constraints I want to impose on what happens to the portfolio. Then, we use Solver in Excel, or something equivalent, to find the best answer. This sounds easy, but there are some big challenges. It doesn't tell you how to do every trade, but it might tell you what a good hedge portfolio should look like. We look for a stable portfolio that we don't have to trade very often that is able to produce the option payoffs at a reasonable price. There are lots of other things that you could do.

We have a highly path-dependent option, so you have to keep track of every single way you can think of to go through a lattice. So, how you've defined the state space makes a big difference in how many different things we have to keep track of. Do we need to keep track of 1,000 different market paths or is 100 enough? Try to limit that and the number of states that you keep track of to make it computationally feasible. Similarly, make sure you're including enough options to come up with reasonable answers. You can think of an unlimited number of options to include in the model and we've struggled to keep it down to the right number yet not eliminate any good solutions.

From my perspective as portfolio manager, it's also very important to me when I use the lattice to price options that the option prices look right against the market. That allows me to calculate an average outcome. This way, when I talk about optimizing the expected cost of hedging or the maximum loss that will happen no more than 5% of the time, my methodology makes some sense. If I just arbitrarily create a set of scenarios equivalent to the New York Seven, there's no way to combine the results. With my methodology, I can.

If the optimization is not convex, you might find an answer that the optimizer thinks is optimal but really isn't. It's locally optimal but it's not globally optimal. I've thought about some approaches to attack that. It definitely is an issue, to be sure, when one is talking about mathematical models. You have to make sure that you believe the answer and that you've met the specifications that will be able to give you the right answer.

One challenge we've had is the need for some serious staffing and other resources to be able to do this well. Insurance companies are increasingly bringing in people who have the expertise to set up this kind of modeling and execute it. However, it's not necessarily the same expertise that people have developed in the past in insurance companies.

Finally, now that I've told you how hard this is, I want to tell you some of the benefits so that you can understand why I'm spending time on it. The optimization benefits are the obvious carrot. If you can save 1% of the value of the premium in hedging costs, that could increase your profitability very substantially on one of these products. If you look at the number of transactions that you have to go through and some of the things that you need to do for a product like ours, it's possible to pick up 2% or 3% on the premium in value. That's the range I like to aim for. We haven't gotten there yet, but that's a big carrot.

In addition, the ability to do stress testing and define a distribution of the outcomes over a certain period of time in your portfolio help communicate your risks to management. I see it dovetailing very nicely with the risk management environment based on value-at-risk or similar methods. It requires you to have internally consistent modeling and to do things in a way that they can be combined statistically.

We haven't started to deliver the results of these optimizations to our senior management yet, but we're definitely pushing in that direction and preparing the ground for it. From a practical standpoint, one of the most difficult things about dynamic hedging is that the business planners ask, "What are you going to spend on options next year? What are you going to do with futures?" I'd like to say, "It depends on what the market does," but can't really do that because they have to tell the analysts something. They have to tell our holding company something as well. At least this allows you to say, "If this happens, this is what we should be doing with the portfolio." You learn about what the right trades might be if the market does various things and that kind of information. Although it's hard to summarize, it gives you a sound basis for business planning. When you're potentially spending a quarter of a billion dollars a year on option premiums, management likes to know where you came up with those numbers.

You also feel much more comfortable managing the portfolio when you can communicate your dynamic hedging strategy and provide a lot of detailed information to senior management to show that this has been well thought out. You've looked at a wide range of scenarios and determined what should be done in each instance. I think that's a great tool.

Mr. Andrew S. Chow: CCM manages about $30–40 billion in assets, most of which are insurance assets. Some assets are held by our parent and some by outside insurance companies. One of the portfolios I manage is CONSECO Annuity Assurance, which used to be called American Life & Casualty. That company is currently the largest underwriter of EIAs in the country. And we currently manage asset programs for four different companies writing EIAs.

Because of time constraints, I'll just talk about issues relating to the asset side on hedging. There are many issues relating to underwriting P&L, but I won't deal with them. There are two main objectives in the hedging business—speed and accuracy. Both relate to what we are trying to do—minimize any possibility of loss from our hedging operation. Bear in mind that you want to sell EIAs to make money. You have different sources of profitability and loss. One is underwriting P&L and another is risk of bond default, which I want to minimize. Then I have some friction between how I hedge the options that I've granted to the policyholder and the options that I buy on the asset side.

Speed and accuracy both lead to better hedge results. Accuracy is pretty easy to understand. If I don't know what I'm trying to hedge against, it's like driving with my eyes shut. The probability that you hit a signpost or mow down a pedestrian is significantly higher. But what about speed? The sooner I get results back that tell me where I am in the hedging process and how close from my targeted hedge position I am, the better. If I get results back on my hedging process that are a week old, it's like driving by looking in the rear-view mirror. Once again, it's not very effective.

What are the obstacles to speed and accuracy? The biggest problem in terms of getting good, accurate hedge results is operational in nature—things like your sales force reporting a different sales number from what actually occurred. It sounds simple, but it's actually the meat of the problem. It has nothing to do with deltas or gammas. I can brainstorm for a long time about my best hedge strategy, but then I get a phone call that the $15 million in sales reported was actually $10 million. That throws you a deep left hook when you're not expecting it.

Backdating is another problem. There are people in sales support whose job it is to make salespeople happy. One of the things they do is slide an application under the table when the salesperson who has missed a deadline by 24 hours wants to get his or her commission check next week. Once you start allowing that to happen for a couple months, people begin to realize that backdating allows them to get a more favorable mark for their customer on where their index start date is. This is a big problem because the S&P index could be down 5% in the course of a week, and if you allow people to backdate, you'll find that people are systematically backdating when the index goes up but not when it goes down. They want to go back to the time period when the index was lower. If your people are doing that and taking out 5%, you're going to get waxed. That will be your entire P&L for a year. Never mind about the delta or gamma, you're just going to get killed. So, these are the practical problems that we have found to be the biggest issues in terms of maintaining a hedged P&L at zero.

Model error is another source of problems. Depending on how complex your option is, this could be an issue. Calculation time is also a problem. This relates to the speed factor. Our options are fairly simple, but they don't have a closed-form solution for determining the market value. The Black-Scholes formula has a closed-form solution. As long as I know what the parameters are, I can solve the equation, and it takes the computer two nanoseconds. If I have a position of 400 options, it takes 800 nanoseconds. Unfortunately, our options are a little more complex. We're valuing something on the order of 400 options on the asset and the liability side, and we want to test those each in five scenarios because we don't want just the market value. We want to calculate the impact of various parameters. We could do it for 20 scenarios, but we have problems even with five. And for each of those options, I need to trace 20,000 paths to get a good level of accuracy.

So, 400 options times five scenarios is 2,000 times 20,000 paths. Now we're up to 40 million paths that I have to trace. And, for each one, I need about 20 steps in terms of price diffusion points. It can be more or less, depending on what month they fall in and how long-dated the options are. Now I have something like 800 million price diffusions to calculate. This is a speed issue. Each one of those price diffusions hits a random number generator and does a number of other calculations, so the total number of instructions I need to give the computer is massive. If I don't get the results back quickly enough, I'm driving through the rear-view mirror. We like to price our portfolio daily, so realistically I have about 22 hours to calculate all the numbers I need. That's the speed issue that I was talking about.

How do we solve these problems? The first thing is to make the process modular. Rather than have one central processing unit (CPU) crunch all the numbers, I'll chop up the process into a number of different pieces. Then I'll get one box with multiple CPUs in it so I can solve the problem more quickly. I want to measure the market value of my hedge position on a daily basis so I can make any changes that are necessary. If I'm only adjusting my hedge weekly, I could have a bad week and, before I know it, I'm off base. Then, in order to communicate our information to our clients, we post it on an internal Web site so people can view it a day after all the calculations are done. When our clients have sales on Tuesday, by Thursday morning they can see how those sales adjusted their risk positions.

All the asset information comes from CCM and our accounting group, and the liability information comes from the individual clients. They provide their sales information on a daily basis. That means that they have to have systems in place to be able to report sales data. Many insurance companies don't know what they sell for three to four days after the sale is made. Clearly, this is unacceptable for this line of business. For a regular interest rate annuity, three to four days of interest rate changes is a one-sigma event, maybe ten basis points. Who cares? That's ten basis points of profitability on a product that probably had 80 basis points. You can make it up in the second year. On this product, however, three or four days of price sigma on the S&P can be 300 basis points. You cannot have a 300-basis point loss on this product and expect it to be profitable. That's why you must have significantly tighter information flow—in terms of backdating policies and reporting of sales—to get good results.

We gather all the information that you would normally think of on an accounting sheet or an asset listing. The interesting stuff is the interaction between the assets and the liabilities. The liability information is presented daily. Most people in companies that are selling EIAs are aware of the difference between daily start dates and weekly start dates. Daily sales create certain issues from a hedging perspective.

The information content is the in-force amount, the participation rate, the number of units, the maturity date, and the strike. Then we aggregate the daily liability information on a weekly basis to compare it to our assets. Now we're starting to get some interesting information. This is what we call the hedge variance report. With it, I can compare the market value of assets with the liability market value.

First, I want to see if my hedge worked. The hedge variance report shows me a decomposition of the hedge variance between various sources of error. The bottom line tells me the net P&L for all my hedge positions. As a percentage of the total amount in force, in general, we've found that it's relatively simple to keep the hedge results variance to within ten basis points. That's a good result. If you do a good job of getting sales reports from your sales force, keep backdating in line, and make sure that the numbers that they report to you are accurate, a variance of ten basis points of profitability since the time of inception—not per annum—is clearly an acceptable number. You are going to see ten basis points of variance on a regular annuity product. For EIAs, this means if you stick to your knitting, you can get tight hedge results that are as good as a regular annuity.

What if somebody tells you that the market value of all the assets is greater than the market value of all the options that you granted? Your first reply should be, "Who cares? It's not a relevant number." Total P&L is determined by a host of other factors. For example, the delta amount for one portfolio shows a net short of 247 S&P units. That's equivalent to two futures contracts. However, it's totally insignificant in comparison with the size of this portfolio, which is $500 million. The gamma (the old textbooks used to call it kappa) number shows a net long of about $238 for a 1% change in the market. Once again, who cares? It's a $500 million portfolio. This demonstrates that the total hedged variance is very insignificant, if properly constructed, relative to the other issues I've talked about, such as reporting errors and backdating. Those are the big ones to keep your eye on if you want to do a good job on the asset side.

What happens to the portfolio when the market value goes up and down? If the S&P 500 goes up 1%, I lose $13,000; if it goes down 1%, I lose $8,000. So, regardless of whether the market goes up or down, I lose money. This is in direct contrast to the kappa or gamma number, which shows that I'm positively convex. This is another good illustration of where the greeks can fall down. After making all these careful calculations the gamma measure failed us.

Mr. Glacy: Why is that?

Mr. Chow: The kappa number was calculated using a shift size up and down of 1%. In contrast, another kappa number was calculated using a 0.50%. So, there's another assumption behind all the assumptions. We also have the vega number for changes in volatilities up and down 1%. The only thing I have to add is a market issue. We estimate that total insurance sales of EIAs have generated approximately 15–25 million units of vega demand for options volatility this year.

What does that mean? Well, it's no secret that the largest seller of vega recently blew up. The company sold vega units in excess of the total demand for vega from all insurance companies in the U.S.—and it's gone. This is a significant issue for EIA writers that will bring a lot of volatility to options pricing. It's going to be a relatively treacherous market.

Mr. Larry M. Gorski: I have a question for Jeff. You talked about a stable hedging portfolio, which seems to be somewhat different from the context of the last speaker. When you say a stable hedging portfolio, are you talking about stable for a day, a week, a month, or a year?

Mr. Lobo: There are various levels of stability. If you have a simple liability option that can be easily purchased in the market, you can buy one option, hold it until the end, and use it to pay off the policyholder. That's the most stable situation. The opposite situation is trading futures all day long, every day, to try to replicate the option. You're probably going to have to trade a half-dozen times a day, at least, to keep your portfolio flat. That's not stable at all, and it costs a lot of money. Futures costs aren't that high, but they add up.

Until about a year ago, we had a hedging strategy that was fairly well-defined insofar as the trades we would make every year. We were turning over our portfolio about once a year, and it was costing a lot because the transaction costs were very high. If I can cut my portfolio turnover from 100% to 50%, I could save millions of dollars per year in transaction costs. That's the level of stability that I'm talking about. I'd like to build a portfolio on day one that will minimize the expected number of transactions that will be required in the future to keep it flat. However, I can't say that I have the best starting portfolio yet.

Mr. Brian Kavanagh: Has anyone on the panel considered the implications of regulation ZZZ as far as meeting hedged-as-required criteria and the possible constraints that would that put on a hedge's performance? If you're not trying to meet those criteria, would you take into consideration fluctuations in reserves in managing the assets?

Mr. Stone: Fluctuations in reserves are a double-edged sword. To the degree that fluctuations in reserves offset the mark-to-market on our hedges, and we have our tracking error under control, we wouldn't be concerned about not being hedged-as-required.

Mr. Chow: Our read on the ZZZ regulation is that it's more of an issue for the longer dated products. Producers of longer dated products have tried to exploit lapse assumptions to allow them to hedge less than the full amount of sales. For a short-dated product, we're adjusting lapses dynamically on a weekly basis, so it doesn't affect us. It's not a factor for short-dated and annual-reset products.

Mr. Lobo: As a seller of long-dated products that does a lot of dynamic trading to hedge its portfolio, we believe that it would be uneconomic for us to meet the hedged-as-required criteria for our portfolio of liabilities. We're going to have to show the mark-to-market in *Financial Accounting Standard (FAS) 133* anyway, so we'll live with potential fluctuations in our reserves. We don't believe they'll be that bad because we're going to be able to meet the test to use our reserving methodology, despite not being hedged-as-required.

Mr. Stone: Following a Type 2 method by having your reserves and hedges mark-to-market is a very good reporting mechanism because your hedging results should flow right through your statutory income and, under *FAS 133*, flow right through your GAAP income. They will be quite apparent. It's very important that external audiences understand the magnitude of the risks a firm is taking by engaging in this sort of an activity. The ZZZ Type 2 method in *FAS 133* will aid in that kind of transparency.