RECORD, Volume 25, No. 1*

Atlanta Spring Meeting
May 24–25, 1999

Session 73PD
Does Anyone Here Speak Greek? Hedging Your Equity-Indexed Products

Track: Investment
Key Words: Investment

Moderator: ANSON J. (JAY) GLACY, JR.
Panelists: THOMAS BAUER†
ANSON J. (JAY) GLACY, JR.
KEVIN P. GUCKIAN‡
HENNING HASLE§
Recorder: ANSON J. (JAY) GLACY, JR.

Summary: Equity-indexed products require an asset strategy that matches the equity options embedded in the liability. The panel discusses the following:

- *Two asset strategies*
  1. *Purchasing equity options that match the liability options: Are you truly matched or are there still risks?*
  2. *Dynamic hedging using the mathematics of "the Greeks": Is it self-insuring or are you just kidding yourself?*
- *Accounting issues with these strategies*

Mr. Anson J. (Jay) Glacy Jr: We'll focus on hedging equity-indexed products, more specifically on the use of dynamic hedging techniques. We have a first-rate panel of practitioners assembled to explore the many complicated and challenging issues associated with dynamic hedging. I'm a consulting actuary with Ernst & Young based in Hartford, Connecticut. Our practice focuses on capital markets risk management for the insurance industry. I'll talk very briefly about the liability side of the balance sheet, the so-called Greeks, and what dynamic hedging is. Then I'll turn things over to the experts.

---

†Mr. Bauer, not a member of the sponsoring organizations, is the owner of Zeus Capital in Westport, KY.

‡Mr. Guckian, not a member of the sponsoring organizations, is Senior Manager at Ernst & Young LLP in New York, NY.

§Mr. Hasle, not a member of the sponsoring organizations, is Vice President of SAFECO Asset Management Company in Seattle, WA.

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

First, Henning Hasle holds an MBA from the Wharton School and has 12 years of international investment banking experience dealing with derivatives and other capital markets products. Henning joined SAFECO in Seattle as vice president and head of financial risk management. Henning will discuss a practical example of dynamic hedging in action. Second, Tom Bauer is currently in the process of launching his own money management firm, Zeus Capital. Previously, he was the derivatives portfolio manager at ARM Financial Group and an investment officer at Providian Corporation and Capital Holding Corporation. He has over nine years of experience in dynamic hedging of equity-indexed options, and his focus today will be on the pitfalls of dynamic hedging. Finally, Kevin Guckian is a senior manager in Ernst & Young's New York office specializing in the insurance industry. In addition to audit responsibilities, Kevin provides on-call consulting services to insurance companies, investment banks, and other financial institutions regarding accounting and regulatory issues related to nontraditional reinsurance, insurance products, financial instruments, derivatives, assets, securitization, and other capital markets transactions. Kevin will focus on the accounting aspects of dynamic hedging of indexed products.

Let me start by defining what dynamic hedging is. Essentially, if you understand the traditional concepts of duration and convexity, you'll understand dynamic hedging as well. It entails establishing a position in certain instruments whose market movements parallel those of the liabilities to be hedged. The key theoretical assumption underlying dynamic hedging is the continuous rebalancing of positions. Finally, the key to dynamic hedging in practice is the use of an optimizer to solve for the cheapest-to-procure portfolio of bonds, futures, and options that best meet the liability targets.

For equity-indexed annuities (EIAs), the types of instruments that a dynamic hedger would use include the Standard & Poor's (S&P) 500 futures contracts, S&P 500 index options, and Treasury futures in addition to the bedrock fixed-income position that supports the product guarantees.

Now, one of the lessons that we learned from 1998 market events is not to minimize the importance of vega and theta. If you look in the Hull textbook (*Options, Futures, and other Derivatives*, by John C. Hull, 1998) the author defines theta as a minor "Greek," but for some insurance companies in 1998 it had major implications.

Some of the leading writers of equity-indexed products are adherents of dynamic hedging, so I believe that dynamic hedging's penetration in the insurance industry has grown. There are other companies involved in it, including Allstate, ARM Financial, CONSECO, Jackson National, Keyport, Life of the Southwest, and SAFECO.

Who are these mysterious Greeks that I've mentioned a couple of times already? Let me define the five major Greeks from an equity-indexed liability perspective. The first major Greek is known as delta, which is the sensitivity of the fair value of the EIA liability to changes in the S&P 500. It's akin to the first derivative. On a shorthand basis, people sometimes refer to dynamic hedging as delta hedging.

The second-order sensitivity is called gamma. That's analogous to convexity on the interest-rate side. Gamma captures the curvature of the movements in the value of the equity-indexed liability. As a second-order measure, it's also the first derivative of delta.

Third, vega is an honorary Greek since it's not actually a Greek letter. Vega is the sensitivity measure of the EIA liability with respect to market-implied volatility.

Fourth, one of the so-called minor Greeks is theta, which is the rate of change of the EIA liability with respect to time decay.

Finally, the real analog to duration on the interest-rate side is rho, the sensitivity of the liability with respect to changes in interest rates. My partners on the panel will address each of these Greeks in more detail as we proceed

What tools are necessary in order to compute liability Greeks? First and foremost is a robust and rigorous stochastic platform. Ensure that your modeling of interest rates and movements in the S&P 500 is rigorous and employs a two-factor model to suitably traverse the event space. The correlations between interest rates and equity market movements will be critical here. Further, in our Hartford-based practice we take advantage of some nice variance reduction techniques, like low discrepancy sequences, that are very handy for purposes of accelerating Greek computations. Finally, since we're computing market values and their sensitivities, the appropriate risk-neutral valuation process is required.

Although the experience so far hasn't been conclusive, we believe that policyholder psychology will be different for EIAs than for traditional products, so we've identified two basic modes of policyholder psychology nicknamed naive mode and savvy mode. The naive policyholder is one who buys when the S&P 500 is "high" and sells "low." Conversely, the savvy policyholder gets into the policy when the S&P 500 is low and then attempts to lock in gains when the S&P 500 is high. Those two disparate psychologies will have a dramatic effect on the calculation of liability Greeks, as will the correlation assumption between the movement in interest rates and the movement in the S&P 500, because it's the level of interest rates in many cases that will critically affect the policyholder's decision to depart the contract.

It's also important when calculating liability Greeks not to stop the process at the next interest crediting date or the end of the current index term. In observing option-pricing concepts, it's necessary to follow the liability until it turns into cash, tracking that liability beyond the next interest rate reset date or index term. In doing so, it's also necessary to formulate a renewal crediting strategy that responds appropriately to market dynamics.

Insert 1 contains a couple of what we call "adaptive" crediting formulas. The first formula adjusts renewal participation rates in accordance with changes in the Black-Scholes option costs. The second formula is an improvement that takes into account how changes in the insurance company's earned rate on assets may supply additional crediting latitude. So, again, it's important to track the policy until it turns into cash and not to stop the calculations at the end of the current index term.

INSERT 1
COMPUTING LIABILITY GREEKS

$$PR_t = (BS_0 * PR_0) / BS_t$$

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

There are some other considerations in computing liability Greeks. First, from a capital markets perspective, the EIA can be viewed quite differently. It's really just a long-term string of policyholder puts. In other words, the policyholder has the right at any time to put that contract back to the insurance company if he or she is unhappy with the policy (under naive psychology) or if he or she wants to get out at the "top" of the market (under savvy psychology).

Second, the textbooks define an EIA as a "compounding-notional floating-strike look-back put." And believe it or not, that's what it really is. There is actually a closed-form solution for the option price and some of the Greeks. It's not quite perfect because it doesn't properly take into account some of the policyholder psychology that I've talked about.

Next, in the way of computing considerations, it's important to recognize multiple index terms, the shape of the yield curve, the implied volatility surface (if you have caps on your product), and the impact of the perturbation shift size. The shift size is the shock to initial conditions that is made to calculate sensitivities. And be sure to integrate minimum guarantees into the Greek—not just nonforfeiture minimums, but any base guarantees you might have as well.

In Table 1 I've contrasted two different ways of calculating the Greek for a simplified five-year, point-to-point structure. The proxy Greek reflects a nondynamic, simplified calculation just out to the next index term. The true Greek, on the other hand, reflects all of the dynamic modeling considerations I've discussed.

TABLE 1
PROXY GREEKS VERSUS TRUE LIABILITY GREEK
FIVE-YEAR POINT-TO-POINT STRUCTURE

Greek	Proxy	True
Delta	0.66	1.03
Gamma	0.0006	0.0024
Vega	7.9	11.6
Rho	-0.64	0.53
Theta	-3.23	-23.76

In contrasting the two columns, the proxy column is nondynamic in nature. It doesn't capture policyholder behavior, be it naive or savvy. It doesn't capture insurance company behavior in terms of resets in the participation rate or the cap and so on. If you look at the true column, it brings all of these dynamics into play, most importantly policyholder lapse behavior. So, in hedging, it's important to establish the true liability Greek as your hedging target. For example, note that invoking dynamic policyholder lapse behavior makes the liability much more sensitive to market movements, as we can see by comparing delta and gamma.

Mr. Henning Hasle: I'll talk about our EIA product and some of the hedge decisions we have made at SAFECO. In common terms, I'll discuss what delta hedging really is and then look at a practical example based on our current product. Then I'll discuss why people should consider delta hedging.

The EIA product we have is a seven-year product based on the S&P 500 index. It features a 100% participation rate less a margin, so, if the S&P goes up by 10% and the margin is 5%, we will credit the policyholder 10% minus 5%, or 5%. Our product resets the index each year. That's what they call a "ratchet" product in the marketplace. It's nice for us because that means it's actually a string of one-year options. And, lastly, we have simple interest crediting, which means we just credit interest at the end of the term and don't compound interest.

Let's be a little more precise about what we have been doing. For us, the option amount for 1999 was approximately $480 million of notional amount of call options. As I said , because of the ratchet nature of our product, it's really only one-year options that we hedge each time. For 1999, we decided to go to Wall Street and buy the options directly from investment banks like Goldman Sachs, and do that for half the notional amount. That cost us $24.4 million in option premium, which we paid to our bankers. It corresponded to a volatility (they always talk about volatility when they talk about the price of options) of 25.24%. We then take the other half of our liability and run a delta hedge. If the cost of running our delta hedge is less than the $24.4 million, it's better for us to have run the delta hedge. If the cost turns out to be higher than $24.4 million, it would have been better for us to have bought all the options from Wall Street. The interesting thing with delta hedging is that you never know until the option expires whether it would have been better for you to buy the option from Wall Street or to delta-hedge it. It depends on how the markets developed over the year.

What is delta hedging and all these differential equations? Basically, it's hedging options with a portfolio of underlying assets such that the change in the portfolio value equals the change in option value for a small movement in the underlying assets. In our example, we are using S&P 500 futures to hedge our EIA liabilities, so what we want to establish is a portfolio of S&P 500 futures that change in the same amount as our EIA liability for a small change in the S&P 500. Delta hedging works extremely well for small changes in the S&P 500, but not for big changes.

Let's look at an example. Let's assume that the S&P 500 is at 1,260 and we call up our bankers and ask, "What is the option value?" and they say 114.55. If the S&P 500 ticks up by 1 to 1,261 and our banker says, "Oh, now the value of the option is 115.07," what has happened is that the S&P 500 index increased by 1 and the value of the option only increased by 0.52. That is what we call delta. Delta will then be 52% for this option because when the S&P 500 changed by 1, delta changed by 0.52. If we had sold, for example, 10,000 options and bought 5,200 units of the S&P 500, we would be delta-hedged because, for a 1-point movement in the S&P, we would gain 5,200 on our portfolio and lose 5,200 on the options we had written. That is what delta hedging is in actual practice.

Let's look at what happened to the delta hedging in our example. Chart 1 illustrates the volatility, the strike price, the interest rate, and so forth. But let's not focus on that right now. Let's just say that, as the S&P 500 goes up, delta goes up as well. With a strike price of 1,323, if the S&P 500 index went all the way up to 1,900, we would have a delta very close to 1, so we would hold the same amount of S&P units as we had written in the option. It's very likely that the option is going to be exercised, so we want to make sure that, when the S&P 500 goes up by 1, the value of the option probably also goes up by 1 because it is so deep in the money.

If the S&P 500 declined to 500, the value of the option would be 0 whether the S&P 500 was at 500 or 501. So, when the delta is very close to 0 we hold none of the underlying assets. Basically, what delta hedging is all about is adjusting your portfolio of underlying assets such that you always are hedged for small movements in the S&P. So, if you were hedging a written option, you would buy S&P units as the S&P goes up and sell them as it goes down.

Let's take a numerical example to illustrate that the basis of delta hedging for an option that you have written is to "buy high/sell low." I don't know if it's a naive strategy, but that's what you have to do if you want to delta-hedge. Let's take the numerical example again. We know that with the S&P at 1,260 the delta was 52, so if you had sold 10,000 options you would have bought 5,200 units of the S&P 500. Now, let's assume that the S&P 500 went up to 1,270. According to my chart, the delta is going to increase to 53.37. That means, in order to be delta-hedged, we need to buy an extra 137 units of the S&P 500. We can do that at the 1,270 price.

Then, say, the next day the S&P falls back to 1,260 again. Unsurprisingly the delta returns to 52, so we need to sell those 137 units. We would have bought the units at 1,270 and sold them at 1,260, so we would have incurred a cost of buying high and selling low. That is the cost of delta hedging. When I told you earlier that we paid $24 million for our options, if we don't lose more than $24 million by buying high and selling low during that year, we would have been better off delta hedging than buying the option from Wall Street. If we lose more, it would have been smarter buying the option from Wall Street. That's what it's all about.

In our example, we use the S&P 500 index futures contract as a hedging instrument. We rebalance our holdings once per day. That's pretty much what Wall Street does, and it works well in practice.

Let's take a look at our experience in the month of January as shown in Table 2. Our option is set to strike at the end of the year, so on December 31 we bought 432 units of futures, which corresponds to a notional S&P 500 amount of $133.6 million. That was our initial position. The $133.6 million is our delta hedge for the written option of $240 million.

TABLE 2
JANUARY 1999 HEDGING EXPERIENCE

Date	S&P 500 Index	S&P 500 Future Price	S&P 500 Future Trading	S&P 500 Future Contracts	Share Trading in $million
12/31/98	1,229.23	1,237.48	432	432	133.6
1/04/99	1,228.10	1,234.10	-1	431	-0.3
1/05/99	1,244.77	1,253.97	16	447	5.0
1/06/99	1,272.36	1,282.59	27	474	8.7
1/07/99	1,269.73	1,278.20	0	474	0.0
1/08/99	1,275.09	1,285.00	4	478	1.3
1/11/99	1,263.88	1,275.00	-10	468	-3.2
1/12/99	1,239.51	1,249.29	-24	444	-7.5
1/13/99	1,234.40	1,239.00	-5	439	-1.5
1/14/99	1,212.19	1,219.25	-20	419	-6.1
1/15/99	1,243.26	1,249.00	31	450	9.7
1/19/99	1,250.89	1,259.00	8	458	2.5
1/20/99	1,256.56	1,264.00	5	463	1.6
1/21/99	1,235.17	1,240.19	-21	442	-6.5
1/22/99	1,225.19	1,232.75	-10	432	-3.1
1/25/99	1,233.95	1,240.50	9	441	2.8
1/26/99	1,252.31	1,265.00	18	459	5.7
1/27/99	1,243.19	1,249.00	-8	451	-2.5
1/28/99	1,265.37	1,265.70	22	473	7.0
1/29/99	1,279.64	1,284.00	13	486	4.2

We adjust our futures holdings every day as the S&P changes. When the S&P 500 fell from 1,234 to 1,212 we had to sell 20 units. The very next day it went up from 1,212 to 1,243 and we had to buy 31 units. Here is a buy high/sell low kind of cost. When we add up all those costs as we go along we find out what the cost of delta hedging is. That's what vega is about. If the market is very volatile, the cost of buying high and selling low is going to be very high. If it's a very stable market, the cost is going to be low.

I mentioned that we were targeting a 25.24% volatility in the options we bought. Therefore, if the experienced volatility is less than that, we will profit from our delta hedging. If it's more, we will lose. And the way we calculate that volatility is simply the annualized standard deviation of the daily returns, and for this group that's a simple calculation. That's the vega risk.

There is one other risk with delta hedging, the gamma risk, or the risk of market gaps. Delta hedging works perfectly for small movements in the S&P, but if you have a big movement you're in trouble.

Let's assume we have an S&P of 1,200 and an option value of 27. And let's assume that the portfolio is $141 million. If the S&P dropped by 10% from 1,200 to 1,080, the option value would go down from 27 to 14. And the value of our portfolio would have gone down from 141 to 127. The option value went down by 12, so our liability went down by 12. Because our assets went down by 14, we would have lost $1.9 million that very day. Remember, we had $24.4 million in option premium, so we can sustain maybe one but not a lot of these losses.

The other problem is, if it's not a 10% drop but a 20% crash, then it's an $8 million loss. And if it's a 30% crash, then it's a $17 million loss. And that's not the day that top management is going to be happy when you have to tell them that you lost all this money. But that's the gamma loss. And the approximation is 0.5 times gamma times the change in the S&P 500 Index squared; that's the way you approximate the loss. That's the gap loss, so you're really in trouble if the market crashes. You're in similar trouble if the market goes up, but the markets have more of a tendency to go down by 20% than up by 20%. That's the gamma loss.

Let's review the facts of delta hedging. We do not hedge volatility. We experience volatility by buying high and selling low. So, if you have a very volatile market, a delta hedge is going to be expensive. If you have a stable market, a delta hedge is going to be cheap. And, if the market goes down by .5% each day, it doesn't matter. It costs nothing. It's the big moves that cost.

Delta hedging is not efficient in gapping situations. I showed you a 10% gap where we lose $1.9 million. If it drops 1% it's only $30,000. So 1% is nothing. It's the 10%, the 20%, and the 30% drops that really hit you. This kind of hedging is generally cheaper than buying options because you take the vega risk and the gamma risk. I do believe it's a very powerful technique that can be used to do a lot of things. As Jay mentioned, delta hedging is similar to duration hedging and the gamma is very similar to convexity, so you people know these terms.

Why, with all these risks, would we do delta hedging? The reason is that the experienced volatility, which is the bumping up and down in the S&P 500 that we have seen, is generally lower than the implied volatility, which is the volatility derived from the option prices. So we have some cost savings and maybe some diversification benefit. Let's compare experienced and implied volatility.

Not surprisingly, the experienced volatility was pretty high around the stock market crash of 1987. On average, it's in the 13–14% range. However, option prices are currently between 25–30%, which is high. They were all the way up to 40% in October when we had our little minicrash. One of the big debates we will discuss regarding *Financial Accounting Standard (FAS) 133* is its requirement that you price your options according to these very volatile markets, which is the price of options.

When we run our delta hedging, the spread between the two would be our expected savings, and that's why people like to do delta hedging. That's why Wall Street wants to do it for you. Its mantra is, "Don't do it yourself. We want to make the spread, not you."

Let's look at the example I used before, which is the $240 million notional amount of options we were hedging. If I look at the first quarter, the S&P 500 went from 1,229 to 1,286. Buy high and sell low—our kind of strategy. We spent $2.2 million doing that, but we have to do it continually because we have to readjust our hedge; so, we are hedged all the time. That cost us $2.2 million. We had a profit of $4.9 million in our general holding of S&P 500 futures because the S&P 500 went up. We also earned some interest on the $24 million we had in premium, so our total trading profit was about $3 million.

Let's look at the option that we actually bought from Wall Street to see what happened to its value during the quarter. The S&P 500 level went up, so the value of the option actually went up from the $24.4 million we paid for it to $26.3 million. We had an increase in option value of $1.9 million, so we saved $2.9 million or had a gain of $2.9 million on our delta hedge. The option we were tracking lost about $1.9 million, so we had a profit, by delta hedging, of $1 million for the quarter. That's because the experienced volatility was lower than the implied volatility. This is how the delta hedge turned out for the quarter. The experienced volatility was about 21% while the implied volatility was at 25%.

So, what is delta hedging? It's simply hedging a portfolio of options with a portfolio of underlying assets. We went through an example where we used S&P 500 index futures, which have very low transaction costs, and daily rebalancing. Why do delta hedging? Because it results in cost savings and because I believe personally that the gamma and vega risks might represent a good diversification bet for an insurance company. It's a different risk than we normally take.

Mr. Thomas Bauer: I'm going to rain on Henning's parade a little bit as I talk about the dark side of dynamic hedging.

I have been managing a dynamic-hedging portfolio for nine years, so I've been through both the good and the bad times and seen when it works and when it doesn't. Hopefully, I can share with you some of the hidden risks of dynamic hedging, which aren't necessarily apparent when you first start thinking about it.

How many of you all have ever bought an option with your personal account? About half of you. There are lots of these spreadsheets and option price calculators out there, where you just input some assumptions and a price is magically spit out for that option. You enter the underlying stock price, strike price, whether you want a put or a call, some sort of interest-rate assumption, what the dividend yield is, how long the option term is, and, in EIA land, a participation rate. And, poof, out comes a price that looks definitive. It also spits out the Greek like magic—the sensitivities of the price to changes in rates, the volatilities, and the underlying stock. What I'm going to try to share with you is why those answers aren't necessarily definitive and, in fact, may be a lot fuzzier.

When we talk about dynamic hedging, Jay has done a nice job of describing that the liability that is being written is really a call option. Henning talked about hedging that call option that has been sold to individual policyholders with dynamic hedging, or a specific variety called delta hedging. Interestingly enough, dynamic hedging really gets back to the basics of option pricing that Fischer Black, Myron Scholes, and Robert Merton back in the early 1970s first theorized about. Their theory essentially was that the price of an option is nothing more than the cost of replicating it. That sounds easy. So, if the cost of replicating an option is $20, then the fair price of that option is $20. And by some magical rebalancing, your cost of replicating the option is $20; therefore, its fair price is $20. So, dynamic hedging is nothing more than option replication—trying to replicate an option using cash, stock, and typically some sort of interest-rate hedge.

As befits academics, certain assumptions were made about the way the world works in coming up with this replicated option price—this $20. It turns out that there are really no constraints to when, how, and at what price you can replicate this option that you can rebalance. There is no limit to how dynamic you can be.

As someone who's worked in the real world like all of you, I know that those constraints are not realistic at all. It is not the real world. You might think, "Big deal, so there are some constraints; so maybe that price of $20 is off by a few pennies." It turns out it can be off a lot more than that. Indeed, even if you are right about what future volatility is going to be, even if you are right about how the stock price behaves in the future, and even if you are right about interest rates, you can be dead wrong about what the cost is going to be—dead wrong in a way that benefits you or dead wrong in a way that hurts you. You can either over- or underestimate the price.

How can that be? Let's talk some more about these constraints and what their real world analogues are. The first constraint that Black, Scholes, and Merton talk about in their unconstrained world is that it's a continuous time world where you can rebalance instantaneously at every instant. The second constraint is that when you rebalance there are no transaction costs involved. Your brokers work for free; there is no bid-offer spread and no taxes to be paid. A final key assumption is that, when you rebalance and you're buying and selling these instruments (like the S&P 500 futures in Henning's example), you're always buying them at a fair price.

To give you a little bit of background on how futures are priced, futures prices track closely the underlying stock S&P 500 price. If the S&P 500 is at 1,260 as in Henning's example, there is always a fair value associated with a futures contract when the S&P 500 is at 1,260. What Black, Scholes, and Merton assume is that you can always buy or sell at that price. In the real world, when Henning is pulling the trigger on his rebalancing, he may be paying more or less than that fair price. Very rarely are you getting or paying the so-called fair price.

The real world is a discrete world. The S&P 500 doesn't move instantaneously by little nicks to the next level. It moves in points or fractions of a point. Sometimes it moves by tens of points, so you don't have the chance to rebalance continuously along the way. There are transaction costs that brokers charge. There is a bid-offer spread, and I just mentioned that, when rebalancing, you are buying and selling mispriced instruments.

What are the hidden risks posed by working in the real world as a dynamic hedger? The first one is that there is a fundamental trade-off between transaction costs and what I would call faithful replication. The more faithfully you want to track your liability and replicate the option, the more it's going to cost you because you're rebalancing more often. You're paying more transaction costs and buying and selling more mispriced instruments, and those costs add up. On the flip side, if you want to mitigate transaction costs, you can do that by rebalancing less frequently. But, if you do that, then your liability might get more out of whack with what your hedge portfolio looks like. How do you manage that trade-off? How can you minimize both transaction costs and the tracking error? How can you maximize faithful replication while also minimizing transaction costs?

The second risk is what I call liquidity holes. Imagine you are all dynamic hedgers and that, in order to execute your order, you have to go out that door and get to one telephone. Let's say the S&P 500 is moving fairly quickly and you need to rebalance. Can you all fit through that door and get to that telephone at the same time? Probably not. It's akin to being in a movie theater when someone yells "fire." Can you all get out the door and execute your rebalancing trade when you have to?

An extreme example of this is when the markets actually shut down and you can't rebalance at all because S&P 500 futures are not trading. That's where Henning's gap risk really can bite you. That's what I call hyperdiscontinuity in the market.

The thing about dynamic hedging is, when you want to rebalance most, the liquidity is not there. You can't buy or sell futures contracts. When the market is chasing or when the market is gapping higher, you can't do anything. Either the markets are closed or your hedging instruments are priced very unfairly. On the flip side, when you least need to rebalance, when things are just ticking along real quietly, liquidity is there. So when you need it, it's not there and when you don't need it, it's there. That's what I call a liquidity hole.

Third, with dynamic hedging every option you're trying to replicate turns into a path-dependent option. This is a little subtlety on Henning's point about the alternative of buying an option at a volatility of 25.2% versus rebalancing dynamically at a volatility of less than that. You would make money. What I will offer is that there are certain paths that the S&P 500 takes that still have a volatility of less than 25.2% where you will lose. The cost will be higher, so not every path is created equal. There are many paths that have an average volatility of less than 25.2% that will have a cost of more than that implied by a volatility of 25.2%. Just run some simulations to see. You can have a number of different paths with the same mean return and standard deviation and the results will be vastly different. So beware of the average. Would you cross a river that is, on average, four feet deep? It depends on the river. If it has a flat bottom, yes. If it has a shallow end and a deep end, you probably wouldn't. But you wouldn't know that just from hearing that the river is, on average, four feet deep. If the average experience volatility over the next five years is less than 25.4%, you might not save any money.

The fourth thing is something called bleed and shadow. Delta changes for the oddest reasons. Delta will change just by today becoming tomorrow. It will change because volatility goes from one level to the next. Delta bleeds. You will incur transaction costs and mispriced instruments because of this so-called bleed in your Greek. Some of the Wall Street firms call bleeding with respect to time "shadow delta." It's a subtlety that you need to keep in mind because, over time, that bleed adds up with these elongated liabilities.

What are the implications of these risks? I mentioned that we had this trade-off between transaction costs and faithful replication. You're going to have to figure out a rule about when you are going to rebalance. Hopefully, you'll figure out a rule that optimizes or minimizes tracking error and transaction costs. Henning and most of Wall Street adopt what I call a time-based rebalancing rule in which they rebalance once a day, once a week, or once an hour. I have always practiced a delta-based rebalancing rule that rebalances only when the delta of assets versus the delta of liabilities gets out of whack by more than a certain amount. Because the market tends to zigzag, a time-based rebalancing rule overhedges, so it is more costly than it needs to be.

Consider Henning's example. He had the S&P move from 1,260 to 1,270 and back to 1,260, thereby incurring some transaction costs (buy low, buy high, sell low) of 1,370 units or dollars. In a delta-based strategy where you would only rebalance if delta moves say from 0.52% to 0.56%, you wouldn't have done anything, so you would not have incurred those transaction costs.

The second thing is that the cost of a dynamically replicated option is uncertain. There is no answer even if you're right about average experienced volatility, interest rates, and dividends. Even if we do our jobs perfectly, we still can be wrong about the cost because of path dependence, the frequency of rebalancing, how mispriced the hedging instruments are, and how many liquidity holes we might step into. We just don't know, so there is fuzziness even if we're right. Now imagine how much fuzziness there is if we're wrong, if volatility turns out to be a lot higher or if rates are a lot different from what we expected.

In my experience, just to quantify this, I think you can be wrong by as much as 20–30% of the option cost because of this fuzziness, and that's real money.

Greek doesn't tell the whole story. I didn't really talk much about delta, gamma, rho, and vega. I talked about the real-world implementation, which wasn't contemplated by Black, Scholes, and Merton or by any of the Greeks. My concerns include liquidity holes, a theater with one door, and a river that's, on average, four-feet deep. These don't show up in any of the option-pricing literature or any of the models, but they're a very real part of a dynamic hedger's job. Dynamic hedging is not at all a tidy science that can be performed by a robot. It's a bit of an art, so you need someone at the helm who has experience in trading these things in the real world.

Here is something to think about with respect to your areas of expertise: Dynamic hedging is a higher-risk, but probably higher-return, activity than simply buying options. What that probably translates into is that there's more capital required by this activity than by passive hedging. There's a lot more that can go wrong here than a counterparty simply not paying on an option you bought from him or her.

The second thing to consider concerns cash-flow testing: How do you know whether you're going to have enough money to pay off your liabilities if these paths really do cause this fuzziness?

Third, when you're pricing your liabilities, even if you haven't sold a penny of this product and you want to put on a dynamic hedging strategy, what price do you use for the option? How much fuzziness do you assign to it? What sort of margin should you expect? What sort of return on capital should you expect? These aren't easy questions to answer.

Hopefully, you all now understand better some of the hidden risks of dynamic hedging, which can have very real economic consequences.

Mr. Kevin P. Guckian: I'm going to spend some time talking about the accounting implications of some of these hedging strategies, especially focusing on *FAS No. 133* considerations. *FAS No. 133* will dramatically change the way these products are accounted for and, while the accounting is pretty straightforward, I believe it's going to be somewhat controversial too.

Before we get into exactly how you account for EIAs in the new framework, let me talk about some of the key concepts of *FAS No. 133* and where the accounting is going. First of all, there are a number of key concepts that the FASB came up with back in 1992 when they started this project, and those key concepts stayed the same throughout the whole project. First, derivatives are, in fact, assets and liabilities. They should be on the balance sheet, so no more off-balance sheet treatment for swaps and other instruments. Second, fair value in the FASB's mind is the appropriate measurement for financial instruments and, in fact, is the only relevant measure for derivatives. Third, only assets and liabilities should be on the balance sheet, so the notion of deferred gains and losses is not within the conceptual framework and, therefore, should not be permitted under the new accounting model. And fourth, there should be some sort of special accounting provided for certain qualifying transactions. This is the old hedge accounting.

The definition of a derivative, according to the FASB, is a specific definition in *FAS No. 133*. When the FASB started the project, it thought about just identifying certain contracts that would be derivatives, such as swaps, forwards, and futures, just as it had done under *FAS No. 119*. But it soon realized that that list would quickly become obsolete and came up with a definition that relied on distinguishing characteristics. The derivative contract must have an underlying (like an interest rate or a security price) and a notional amount—a number of units, for example, that the underlying is applied to in order to determine the settlement values. A derivative has minimal initial investment that's not equal to the notional amount. And, a derivative provides for net settlement. You can go out into the market and settle the contract net.

All derivatives are going to be recorded on the balance sheet at fair value. *FAS No. 133* provides for special accounting for hedges but, because hedge accounting is elective and relies on management's intent, it should be limited to transactions that meet some reasonable criteria that are somewhat rigid.

Derivatives that are not in a hedge transaction, which are not hedging something else, will have changes in fair value that run through the income statement as they occur. For each period, you're going to mark derivatives to market and run them through earnings.

While a contract may not meet the definition of a derivative, it may contain embedded derivatives, which are items that affect the settlement in a manner that is very similar to a derivative contract. An example of these might be a debt security that's linked to some sort of commodity or to some equity return. Convertible securities may contain embedded derivatives plus certain securities that contain caps and floors or collars that will also contain embedded derivatives.

What do we do? In a big change to current practice, we're going to bifurcate the embedded derivative from this hybrid instrument, which is the combination of the host contract and the embedded derivative. If we can reliably identify that derivative, we will bifurcate it and account for it in accordance with *FAS No. 133*. For the host contract, we apply current GAAP accounting guidance. The embedded derivative that we carve out of the hybrid instrument is eligible to be designated as a hedge.

The general rule is to bifurcate, but there's an exception to that. If this feature of the contract is clearly and closely related to the host contract, then we don't need to bifurcate. What this means is, if we had a bond that had a return that's tied to prices of gold, for example, the return is not clearly and closely related to this notion of bond so you need to bifurcate it. In contrast, if the return is tied to movements in the London Inter Bank Offered Rate (LIBOR), for example, because LIBOR is an interest-rate index that is clearly and closely related to a bond, that feature would not need to be bifurcated. However, there is an exception to the exception. If the contract has an exotic feature that would not allow the holder to recover his or her investment or has a return that's at least twice what the initial return was and is at least twice what the market rate is, then the exception doesn't apply.

What does this mean for insurance companies? Well, some investments that they never imagined contained derivatives (and for which they never had to worry about derivatives accounting) are now subject to the accounting standard. Examples are convertible debentures and certain structured notes. In addition, certain insurance products, such as EIAs, and certain variable products that contain guarantees will be subject to bifurcation. For property and casualty companies that write contracts that combine traditional protections with foreign exchange exposures, these contracts contain embedded derivatives. And some of the catastrophe bonds that we hear about will also contain embedded derivatives.

There is one thing I should mention. One of the hedging criteria that is important for us to understand is that, if the hedged item is currently accounted for at fair value with changes in the fair value running through the income statement, then it is not eligible to be designated as a hedge. This will be important when we think about EIAs because they are considered to contain an embedded derivative—a written call option on the S&P 500.

If you have a hybrid instrument with an embedded derivative that is not hedging something, then it's accounted for at fair value with changes running through the income statement. This means that it is not eligible to be designated as a hedged item. The result is that we're going to carry the embedded derivative at fair value, with changes running through the income statement. We're going to carry the hedging derivative, also at fair value, with changes running through the income statement and hope they offset each other. But, as you've heard from Henning, the problem with delta hedging is that this won't always be the case. So there might be a lot of noise under a delta hedging strategy. This accounting model, by the way, is used regardless of what our hedging strategy is. If we're using delta hedging or if we're using some sort of notional hedge, this accounting model is the model you use.

What's our initial carrying value? When we sell an EIA contract, we bifurcate. The first thing we do is come up with the fair value of this embedded derivative. The remaining value is what's allocated to the host contract. So, if we sell an annuity at 100 and carve out the embedded derivative, and the embedded derivative's fair value is 40 based on today's market conditions, then the host contract's value is 60. Two weeks later, if we sell a contract and market volatility is now higher, such that the value of the embedded derivative is now 60, the new host contract has a value of 40. That means we're going to take our host contract and ultimately accrete that up to our guaranteed value. And because we're starting at a lower amount or a higher amount depending on the volatility of the option value, we're going to have a different interest cost on our income statement purely because of the change in the volatility of the market.

For subsequent carrying values of the equity-indexed liability, we'll mark the embedded derivative to market at fair value. Now, *FAS No. 97* requires that a liability for an annuity contract must be equal to account value. There was a lot of discussion about this when *FAS No. 133* was issued and being deliberated. Isn't *FAS No. 97* at odds with *FAS No. 133* because, as I mark the embedded derivative to market, isn't it possible that my liability could go below account value?

The FASB says *FAS No. 133* and *FAS No. 97* are not at odds. We must carry our liability in an amount that is at least equal to the account value. So, the sum of the host contract and the embedded derivative cannot be less than the account value. In any situation where changes in market conditions cause the derivative's fair value and the host contract combined to be less than account value, we must adjust the liability up to account value. Those changes will affect earnings; they'll run through the income statement.

When I prepared this presentation, the effective date for implementation was the year 2000 for calendar-year companies. But last week the FASB voted to delay implementation of *FAS No. 133* for one year, providing calendar-year companies a one-year window before adoption in the year 2001.

Mr. Scott D. Houghton: One of your examples, Jay, had a proxy delta of 0.66 and a true delta of 1.03. You did say the difference is due mainly to policyholder lapse behavior. I wonder if you could elaborate on that. It seems like the policyholder has the option to surrender and that the option is at 90% of the account value growing at 3% interest. That doesn't seem that valuable of an option. I wonder if you could elaborate as to what really drives that difference.

Mr. Glacy: That's a good question, Scott. The purpose of that example was to contrast the difference between a simplified proxy calculation of the Greek and a more robust calculation that reflects some of the important dynamic elements. The proxy was purposely set up to be very simplistic. It was a five-year European option without any dynamic behavior either on the part of the policyholder or the insurance company. And the Greek it produced is the type of Greek that you would generate if you were calculating Greek in a static framework. Then we started to turn on the dynamic elements. We incorporated dynamic lapsation (but I don't recall whether it was naive or savvy) and extended the calculation out either three or four index terms, which then allows the writing insurance company to reset the participation rate at those points. That accounts for the dramatic difference, if you noticed, in theta that goes from 3.23 units to 23.76 units.

The change in the delta you mentioned really reflects an increase in the policy's overall sensitivity to how changes in the market today affect the policy's ultimate tendency to turn into cash and for how much. The challenge is to confront these dynamic elements that are rooted in psychology and incorporate them into a very complicated option-pricing assignment.

Mr. Boris Brizeli: My question is for Tom. Since you have been doing dynamic hedging apparently on the Street, in your experience, how long does it take one to become a dynamic hedger? Let's say you took somebody into apprenticeship. What would you consider an experienced dynamic hedger?

Mr. Bauer: It's one of those things where you can never have enough experience. I think the key is to have been through some sort of dramatic moment, and just one is enough. You need to have experienced a liquidity hole, for example. Going through at least one episode where you can't rebalance is all the experience you really need. I can't predict when the next one is going to come, although it's coming soon.

Mr. Brizeli: Let me put the question differently. One book I read on dynamic hedging says that approximate apprenticeship time for a dynamic hedger running a vanilla book is about three years. Would you agree with that statement?

Mr. Bauer: Yes, that's probably a good time frame. I got into this in 1990. One way to get experience is to try it on a small portfolio. And make sure it's real money. I'm not convinced that paper portfolios, where you go through the motions, give you the sort of experience you need. I think you need real money on the line. That will force you to make the decisions on rebalancing rules. That will force you to confront situations when your broker doesn't pick up the phone and things like that.

Mr. Hasle: I have something to add. I think it's also very important *where* you get the experience. You might get it if you sit on Morgan Stanley's trading floor for a year. If you sit in an insurance company on a trading floor, you might sit there a very, very long time without understanding it. It's very important when you get that experience to have somebody there who knows what's going on. One of the problems I've seen many times is that these formulas are pretty easy. You can buy a little spreadsheet and get the deltas and vegas and gammas out of them, but what do they really mean? When you are in the heat of a battle, what do you do? If futures are "limit-down" and the market is crashing, what do you do? People who have been through it on trading floors on Wall Street have the experience. They have been there before. Insurance companies are not used to this kind of thing because they just sit on their portfolios. So, I would not recommend to anybody just to take a textbook and say, "This looks easy. Let me try to do it in my insurance company. Give me $100 million and let me run a delta hedge." I would strongly advise against that.

Mr. Francis P. Sabatini: I have a question for Tom Bauer and Henning Hasle. As an experienced delta hedger, one of the things that I've learned is that shift size is pretty important in terms of calculating your hedging targets. For example, how much do you perturb the markets in calculating your delta? Is it an infinitesimal shift or is it a 10% move in the market? What do you think the appropriate shift size is and how does it relate to the whole transaction cost issue?

Mr. Bauer: A bigger question is, which model do you pick? Ernst & Young has a very robust liability option-pricing model. There are third-party software firms that have numerically based Monte Carlo or binomial-tree-type models. They don't tend to take into account the psychological elements that Jay talked about with respect to policyholder behavior and insurance company behavior. And each of those off-the-shelf models that you buy or develop yourself must make certain assumptions about how their Greeks are calculated. How much do you perturb the tree to generate the Greeks? I have never gotten into the detail of that. But, mainly, because the sorts of options that I've been dynamically hedging have minimal liquidity, how much you perturb the tree doesn't really have a big impact on the Greeks.

Mr. Hasle: At SAFECO, because it's a one-year option and because the Greeks don't change that much, we simply hedge once a day and don't care what the movement is. With regard to the notional amount, we get a weekly update on what the notional amount is and adjust accordingly. In the beginning we adjusted every week, but it was too small. Now we accept $2 million in movement of the notional before we change our hedges. In practice it works pretty well. We have an easy product to hedge because we can use S&P 500 futures and the transactions costs are limited. The real risk we have is the gamma risk because the futures markets close and so does the S&P 500 market when it crashes.

CHART 1
DELTA HEDGING
(x = 1323, vol = 25%, t =1 yr, r = 5%, d = 1.5%)