RECORD, Volume 23, No. 2*

Montreal Spring Meeting
June 19–20, 1997

Session 22PD
Equity-Indexed Products—Design, Pricing, and Hedging (Part I)

Track: Financial Reporting/Investment/Product Development
Key words: Product Development, Equity-Indexed Annuities

Moderator: LILIA M. SHAM
Panelists: CHARLENE M. BARNES
ALAN K. RYDER
Recorder: LILIA M. SHAM

Summary: Interest in equity-indexed products has been mushrooming in both the U.S. and Canada. Equity-indexed annuities (EIAs) have been called the most significant individual product development since universal life. But for many companies, the "how to" problem has been lagging behind the "can do" attitude. To assist actuaries with the "how to," this two-part session addressed the following issues:

Customer and Distribution Perspectives
- *history/variations of design*
- *sales results*
- *target market and customer needs*
- *life versus annuity products*
- *SEC issues*
- *policy forms*
- *market conduct issues*

Pricing and Design Perspectives
- *pricing techniques*
- *design variables*
- *pricing/marketing trade-offs*
- *sensitivity of pricing to market conditions*

Funding and Hedging Perspectives

- *terminal and intermediate hedging objectives*
- *modeling and assessment of risk*
- *C–3 risk of various designs*
- *funding approaches—pros and cons*
- *practical issues in funding*

Ms. Lilia M. Sham: We have a distinguished panel to talk about these issues. We have Alan Ryder and Charlene Barnes. This session will be on product design considerations and hedging issues of equity-indexed products. The second session will be on reinsurance issues and also tax considerations of investments related to EIAs.

Let me give you an introduction of the first two members of the panel, Alan Ryder and Charlene Barnes. Alan is the president of Insource, Limited, a consulting firm specializing in marketing development projects, which include product design projects for both U.S. and Canadian companies. Charlene Barnes is with Union Bank of Switzerland. She works at the equity derivatives desk and is involved in insurance risk management. She has been helping insurance companies in the U.S. develop EIAs in the last year and a half. Alan will discuss the general pricing considerations of equity-indexed products. Charlene will focus on the hedging aspects of these products. But before they start, I would like to give you a market overview of these products.

EIAs have been sold in the U.S. for more than four years. In Canada this product was only introduced last fall and it's mainly sold by banks. Insurance companies entered this market at the beginning of this year. In the U.S., EIAs became a hot product in 1995. It is estimated that about 35–40 companies currently have a product in the marketplace. Perhaps I can do a quick survey of this room. Which of you, or the companies that you work for, has a product in the marketplace already now? That's quite a lot. Which of you, if you can disclose that, is planning on having a product? That's good. Of course, the number of products available in the marketplace is more than that, because some companies actually have more than one product in the marketplace. We have witnessed a proliferation of designs in these products. The total sales in 1996 was estimated to be about $1.5–2 billion, primarily dominated by a few market leaders. It is important to know that many of the products were launched in 1996. So a lot of the products only got partial-year exposure. Therefore, 1997 is a critical year for these companies to measure the success of their products. The sales outlook for 1997 has been bullish and you will hear some of the sales figures—anywhere between $4 and $12 billion, depending on the source, so this will be a very important product for the annuities market in the U.S.

As an extension to EIAs, we see equity-indexed life products and equity-indexed immediate annuity products coming into the marketplace. We first saw equity-indexed life products introduced last summer, and a handful of companies are rolling out their products this year. We have also seen at least one equity-indexed immediate annuity product in the marketplace, and we are expecting to see a few more coming out this year.

I'm going to ask Alan to talk about the general design issues and some of the hedging issues.

Mr. Alan K. Ryder: My task is to go over equity-indexed products from a macro perspective. I'm going to talk about the fundamental issues in design and also talk about the fundamental issues in hedging. We'll get a little deeper into those issues as these sessions progress.

My agenda is relatively simple. I'm going to take an initial look at hedging and pricing. We will look at some of the main product design parameters. I'm going to look at hedging at the maturity value. There will be some discussion about what to do about intermediate values or dealing with surrenders. We'll then go back and look at the basic approaches you can take to hedging. Finally, I'll close with a few remarks about counterparty liquidity and risk-based capital (RBC) issues in the hedging process.

I'd like to start with a look at the end. The easiest way to think of funding an equity-indexed product is to say—what do I do to fund the guarantee, and what do I do to fund the equity linkage? The funding of the guarantee is typically thought of as zero-coupon bond for the amount of the guarantee. Then you go out and buy some sort of call option to get the equity linkage you want. That simple view assumes that nobody dies and nobody surrenders, which is the premise of a lot of what I will talk about. But looking at it that way gives you an analytic framework for designing a product for pricing it and for assessing your risk exposure.

Chart 1 graphically shows what that looks like. One buys zero, it doesn't matter what the market does, it will pay off for 100% of whatever. One buys a call option, and when the market goes up, it has some value, and when the market goes down, it has no value. Very simply one prices this by doing some modeling work to determine what expenses and profitability are appropriate for the product. Then one goes out and prices the guarantee with a zero. Whatever is left over is what is available to buy equity linkage, and you fiddle with that to determine what sort of equity linkage you can actually offer. It looks something like Table 1. These numbers are sort of hypothetical but also sort of real. In this particular example, expenses are 8% of premium, profits are 2%. You can see that the guarantee is the lion's share of the premium at about 70%, and this example leaves you about 20% of premium for equity linkage.

CHART 1
ZEROES/CALL CONSTRUCTION

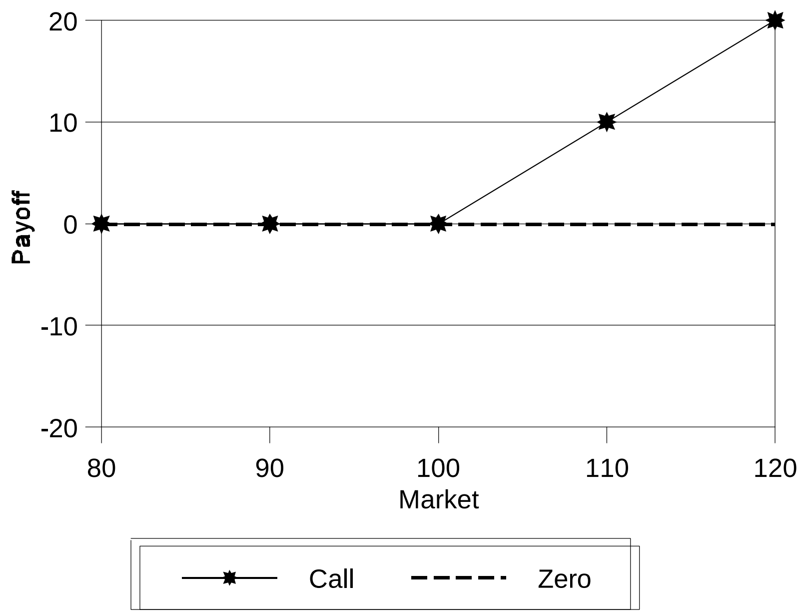

TABLE 1
EIA PREMIUM ALLOCATION

Guarantee	70%
Equity Linkage	20
Expenses	8
Profit	2

The main design parameters you can play with in this product, and there are many others that you can play with as well, including term, guarantee, equity index, participation rates, averaging, lock-ins, resets, caps, and surrender benefits. We're going to look at a lot of these, but not necessarily all of them. We'll dwell on the ones that I consider to be perhaps the most interesting and leave the boring ones behind.

The first one that I want to talk about really is term or the length of the guarantee or the length of the product. Chart 2 shows you the effect of stretching out term on the cost of funding this beast. The cost of zeroes goes down with term, as you would know, and the cost of the options goes up as term extends out. Options' costs don't go up as fast as the zero prices go down so therefore a longer-term product gives you more margin to play with. If you thought that you needed 10% of premium for profit and expenses, this suggests that this product has to be eight or nine years long to make sense.

CHART 2
IMPACT OF TERM ON MARGINS 90 X 1.03^t PRODUCT

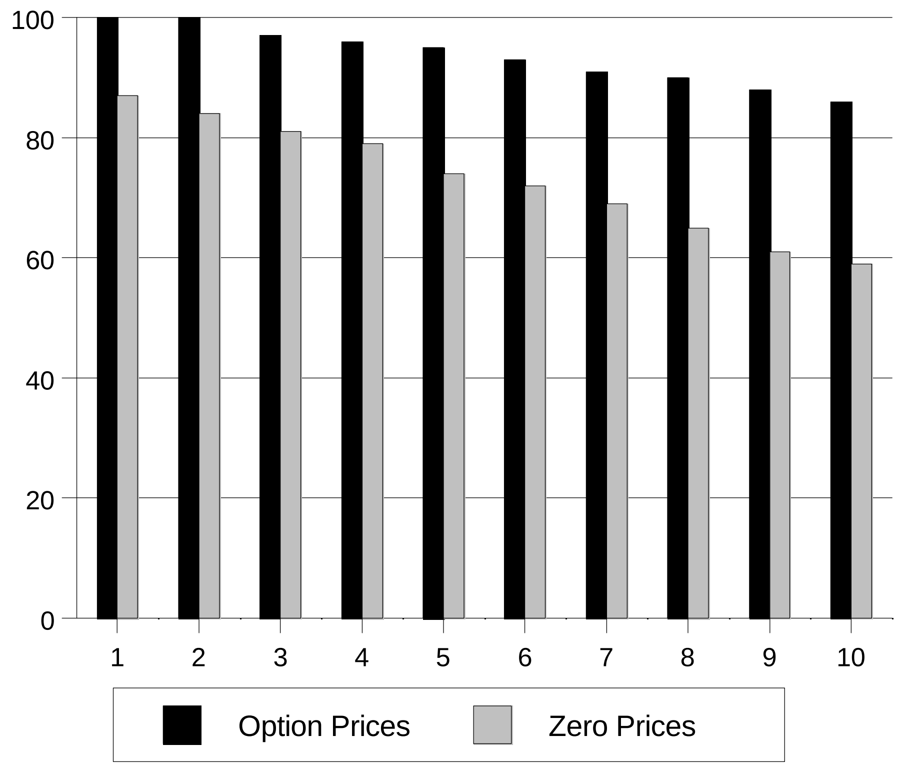

Let's turn to the guarantee (Table 1). The guarantee is very expensive and the issue there, from my point of view, is that whenever you decide to extend the guarantee beyond what the marketplace seems to require, you'd better have a very good reason for doing so. Table 2 shows some representative zero prices, which show you what that's like. The zero prices in the left column are to buy 100, and the zero prices in the right column are to buy a zero for 90 plus 3% a year, which is sort of the most fundamental design one can think of in the U.S. You can see that certainly as you go out, the zero is less expensive, although there's a divergence between these two columns because the 3% guarantee is costing you 3% a year.

TABLE 2
COST OF GUARANTEES

Term	Zero	Zero for 90 x 1.03^t
3 years	86.0	84.6
5 years	74.3	77.6
7 years	63.5	70.3x

Based on recent Canadian Yield Curve

Let's move on now to the issue of what kind of index makes sense. Why is everybody in the U.S. market is using the Standard and Poor's (S&P) 500? I think the fundamental reason is that there are licensing problems with the Dow Jones. From a marketing point of view, the Dow Jones Industrial Average probably is better. It is certainly the best known index in America, probably the world, and it is extremely liquid. Those are two good things. It is not a measure of the total market, which is perhaps not relevant when it comes to marketing, but the fundamental issue is licensing. The S&P 500 sort of scores in second place here. It's well known, it's broader based and it's easy to license—that's why people use it. There are others that you could use.

If you want to use other indexes, what do you need? Any liquid index will do. Liquidity is an issue because the companies that manufacture the derivatives that you want to buy have to go out and actively hedge their positions in the market, and they need to be able to trade stock to do that. If you put them into an illiquid index, you will pay for that or they may not be able to do it at all. In Canada we have a liquidity problem. The principal index that you'd think of using is the TSE (Toronto Stock Exchange) 300, and about 250 of those stocks are not very liquid. So that's why in Canada you would see generally TSE 35, the TSE 35 products being the 35 largest companies in the TSE 300. But in theory, any liquid index will do.

You can choose the ten largest corporations in the U.S. or any of the S&P subindexes. You can use foreign indexes. Now in the U.S. almost universally the S&P 500 is used, and it's a play on the U.S. market. In the U.K. where this product has earlier roots, almost all the products are Financial Times Stock Index products, which is their index. In Canada we have a tiny stock market, and we see the emergence of some foreign indexes products as Canadian investors try to get global and try to get beyond the local market.

You can use foreign indexes with or without currency risk. If you're doing it without currency risk, then it becomes—and this is a very technical thing—kind of an a foreign interest rate play. Foreign interest rates are lower than domestic rates and you can get an exciting product. Right now, Japan is a very exciting place to be indexing.

You cannot use actively managed portfolios. You can't hedge active mismanagement. You can't hedge active trading basically. You can do some surrogates there, but fundamentally there is a theoretical problem with doing that.

I want to digress very briefly before talking about some other pieces of the product. I want to talk about why this product is interesting to consumers. I think on the most fundamental level, this product is interesting to consumers because it's their asymmetrical risk tolerance. They are generally unprepared to wager the principal, but they might be prepared to gamble their gains and that's exactly what this product allows them to do. It secures principal, giving them a variable upside.

An extension of that is once they have their cake, they'd like to eat it, too. Notwithstanding this relatively simple promise of the product, a lot of focus group work has suggested that consumers would likely lock in their gains along the way somehow. This, I think, is a way of extending the whole consumer psychology. I don't want to lose anything, but when I've got it I want to keep it.

We can talk about a number of methods of gains retention. I've listed three: averaging, lock-ins, and resets, and I'm going to cover each of those in turn. You see all these right now in the U.S. market.

All of this is really a definition of what index growth is. All these products pay off the greater of the guarantee or the growth in the index, so now the question is—how is one defining growth? Just define the index; what's growth? You get into averaging, you're defining growth as being the sum ending value divided by the beginning value where the ending value is an average of the index performance over a period of time. That can be a long period of time or a short period of time. There are full-term averaging products in the marketplace, and there are products that average for, say, only the last three months. By the way, theoretically you can average in at the beginning if it has any interest to you at all.

Why bother? Well, it offers the policyholder some crash protection. This particular feature is embedded in every product that I'm aware of in Canada. It's embedded in almost all products in the U.K. and it has some consumer appeal—that's what it's telling you. It says if the market declines at the last minute, then the policyholder somehow is protected from that decline. It's also less expensive to fund, which is perhaps the biggest reason why it's showing up in products. You can see that the difference between the two five-year option prices is quite significant for simply averaging the last year. A rule of thumb, by the way, is if you want to average over the whole term, the option is about half the price if you didn't average at all. If you think about it a bit, it makes some sense.

This is a win/win kind of deal. It's less expensive to fund, which means that you can put more bells and whistles into the product. The consumer likes it because it offers crash protection.

Barrier lock-ins are quite common in the U.K., but I have not seen one in the U.S. They're quite simple. What they say is if the index ever reaches a certain point (and you get to pick that point)—if it ever even reaches it only for one day, then that becomes the new guarantee. The utility of this is that barrier levels can be set at psychological points, and it offers the consumer some opportunity to retain gains. You can build a whole ladder of barriers if you want. Each of them has about the same price because there's a tradeoff of probability of reaching the barrier versus the cost there. The higher up the ladder you go, the less likely you are to get there, but the greater the payoff when you get there.

With this particular style of lock-in, the high water mark is well entrenched in the U.S. market. The index growth is designed to be the highest value of the index at any measurement point along the way. The most common measurement point is annually, but theoretically you can sample any number of times. The cost of this is a function of sampling frequency. As an extreme example, let's look at daily lock-ins. Every day the market goes up, that's credited to your account and every day it goes down, we forgive that loss. You can imagine just intuitively that the cost of that is extremely high. Therefore, one ends up with a very low participation rate or percentage of the growth that is credited to the account. The result is that a product with very low volatility, because the participation rate is going to be low, will give you all the upticks. If you run this through some scenarios you find that this will provide you with actually very low volatility and a low-yielding product. I'd say here that returns are fixed-interest-like.

An annual high lock-in is a more moderate version of what I just described. The cost is moderate, the participation rate is moderate, and you get something that is sort of moderate in the middle there.

The last of the three methods of retaining gains that I want to talk about is recess or ratchets. In this particular style of product, the index growth is defined to be the sum of the product of the periodic increases over each measurement period and again those measurement periods can be short or they can be long. The most obvious one is annual, but they can be anything in between.

Again, if you look at some extremes here, a daily reset will give you more or less the same product as the high water mark daily measurement. An annual reset is what's common in the market. The cost is still high and the participation is relatively low. The sum of style is much less expensive than the product of style in terms of crediting these to the account. Believe it or not, consumers actually prefer that because they can add. They can't multiply, but they can add.

Let's pause and look at cash-out risk of these styles of product. On what I call the point-to-point product, which is a plain vanilla product where the gain is defined as the ending value less the opening value, your cash-out risk as an issuer is the market being down. The market's down and people may want to bail because until the market goes back up, they're under water. The market has to gain something before they're even back to where they were when they started. You have the same thing on an average product except during the averaging period, they're beginning to get some protection. On a barrier lock-in, your exposure is to the market being down before the barrier appears. Once it appears, you're now set at a higher level, presumably materially higher than guarantee, and the consumer is now protected from losses that occur thereafter, so there's less risk of disintermediation. On the high water mark product, the problem is when the market's down before it's up. Once the market is up, then you've done a good job of protecting your consumer, but if the market is in a downturn, you have the same as the point-to-point product.

With a reset product, if the market goes down, reset it down. The consumer is protected from that because it probably has the least disintermediation exposure.

Regarding participation rates, this is my gospel. These are all approximates depending on some assumptions I made. The recessed product will give you, without any other adjustments, the rate and the further averaged product. That is the product where the averaging takes place over the whole term. It is going to give you huge participation rates in what is fundamentally a smaller gain.

Let's talk about the capital asset pricing (CAP) model. There are limitations on the upside, such as 10% per annum or 50% over the whole contract or whatever. There is a marketing issue here. High CAPs might be marketable, but lower CAPs threaten the consumers' conceived value market and they want equity. When you say, "I'll give you the first 10%, but you can't have the rate here," that's a marketing problem. There's also a correcting problem with CAP.

CAP is like the material before it has a pricing impact with this material. A CAP at 20% annual growth doesn't get you very much. A CAP at 10% gets you the low, so you're not going to tradeoff. On the marketing side the consumer might accept 20%, but not 10%.

Now I will just very briefly talk about surrender benefits. What you'd like to do here is offer market value adjustments, but you can't do that. That is desirable, but you can't. There are reasons why many people don't do that. Successes are from a risk management point of view. You do have a minimum nonforfeiture value of 90% premiums per annum. Almost all the products in the U.S. market today are this way.

Your concern is a high-interest-rate, low, soft-market environment (Table 3). That's when the power builder says, I want to get out of this. The other three quadrants here don't really cause you much concern. You can do some scenario testing and you will prove to yourself that the exposure is a relative moderate.

TABLE 3
DISINTERMEDIATION RISK

	Interest Rates Low	**Interest Rates High**
S&P 500 Index Return Low	Moderate Lapses with Slight Gain	High Lapses with Significant Loss
S&P 500 Index Return High	Low Lapses with Significant Gains	Moderate Lapses with Slight Loss

Now I want to switch to hedging approaches. As I said in the beginning, the most fundamental way of looking at this is to buy zeroes and call options, but there are some other alternatives. We'll look at those. First let's look at zeroes and the calls a little bit more completely. You buy zero for the guaranteed amounts, whatever that is. There is a question of how much option to buy and what is the strike price. Both those things are dependent on participation rates and loads. There are a large number of loaded products in the U.K.

In addition, the strike price depends on the guaranteed amount. In particular, if the guaranteed amount is 110% of premium, you need to buy a call that's 10% out of the money because the first 10% is funded. If you do that and you own a lot of zeroes that have low RBC requirements, that's good. They may be causing you yield problems, however. The easiest way to acquire zeroes is to buy Treasuries. Treasuries are not the highest-yielding investments around, obviously. Many companies would like to stretch for yield by taking on some risks, most likely credit risk, and if you want to take on some credit risk here to try to enhance yield, what are you going to do? There are solutions to that. In fact, it turns out that owning a lot of zeroes is perhaps not the best approach at all for this product in the U.S. market.

The second thing you own is a unit derivative. In my example earlier, a derivative goes on your book at 20. Derivatives have high RBC requirements, and this one is a high book value relative to some other derivatives you could consider buying. How would you modify this if you took into consideration death and surrenders or more precisely if you had earlier cash flows? The first thing you might do is replace some of your zeroes, if not all of them, with coupon-bearing bonds. The problem I mentioned before of not getting enough yield begins to go away. You can also add earlier data derivatives depending on the nature of the optionality you've put into the earlier cash flows. You would normally reduce the call option at maturity because if people are surrendering, they're not going to be around at the end to collect on that maturity value.

There are some low-tech alternatives to zeroes in a call. The first is to buy the stock in a put option. That's actually the most obvious hedging solution. The layman might think of that. That's actually a messy solution. It's messy from an RBC point of view. It's messy from all sorts of points of view and it's not highly recommended. A better approach would be to buy futures and floating rate notes. Futures and floaters together give you synthetic stock and it's a much, much tidier thing to do.

Another possibility is to go out and buy floating rate instruments and to engage in a swap. Most swaps that you've heard about are interest for interest, they are typically fixed for floating. In this particular case, we're going to trade floating rate interest for floating equities. Now the currency of swaps would mean that emotionally you have to take that into fixed and then take that back out of fixed into equity, but it's all the same thing. It is certainly possible to engineer the solution. The providers of these assets don't like that one as much because there's now a counterparty exposure that they may get. We'll discuss counterparty exposure momentarily. But this is not a bad solution because swaps have very low RBC requirements.

Finally, one can engage in a structured asset. I think of it as a variable rate bond where the interest rate is determined by some equity linkage. The interesting thing about this is that I think you can categorize, this is not a derivative. I've heard many opinions on that but I believe it's not a derivative. I believe that if you have a problem owning derivatives, that this is a bit of a solution for you.

REPLICATION

Call	=S Stocks – B Zeroes
Call + Zero	=S Stocks + (Z–B) Zeroes
Stocks	=Futures + Floaters

There are some high-tech alternatives to the zeroes and the call approach. There is replication and dynamic hedging. You may recall from one of your actuarial examinations that a call can be synthetically manufactured as a combination of owning stock and borrowing. A call plus a zero is a simple addition to that formula. You can go out and manufacture stocks synthetically by using floaters and futures, and this is how you'd use replication to get to a hedging strategy. What's the problem with doing this? It sort of works on TV, but perhaps not when you get it home. It can be a little difficult to describe to management. In my opinion, this is a good strategy if you have losses of this business to hedge, and it's a poor strategy if you've got very little business to hedge. The cost of getting into this game from a technological point of view and an internal politics point of view is high. This is recommended for those who have lost a hedge. The transaction cost can be high. You're out there trading stock or at least rolling futures all the time. It may not be very practical in general. More sophisticated product designs might be beyond your capabilities. A solution for you might be dynamic hedging.

This is basically the way people like UBS do this. These guys know all about the Greek letters that are associated with options, and they know how to hedge them. I think Charlene will tell us a little bit about Greek letters. These are all the partial derivatives of an option, and there are many of them. This approach to hedging is immunization theory taken to the next level. Immunization theory is all about first and second derivatives of interest rate movements or asset prices with respect to interest rate movements. Here you have more than interest, which is affecting stock prices, and all these Greek letters are some partial derivative for an option.

To do it you need five suitable assets. You solve five equations in five unknowns and you're done. The problem is that it requires sophisticated understanding to know exactly what assets to use and to constantly rebalance this. Because these are partial derivatives, we're ignoring higher moments. Because we ignore those, we drift out of balance. So when companies such as UBS do this, they are rebalancing their options portfolios by the hour, by the minute, constantly, continuously. This requires expertise in software. Most of you don't have that, so don't bother thinking about that.

I have a few closing remarks on some technical aspects of hedging—things you should consider beyond the financial or mathematical impact of the hedging approaches. The first is counterparty. When you engage in one of these transactions, you theoretically have some counterparty exposure. You've turned over some cash and you're expecting some cash back later. That's credit risk or that's counterparty exposure. You have two basic sources of derivatives: exchange-traded and over-the-counter.

Exchange-traded derivatives are highly regulated with daily settlement and minimal counterparty exposure, but they're very limited in scope. Almost everybody in this market are buying over-the-counter derivatives because they want more sophisticated derivatives than the derivatives they can buy on an exchange.

When doing business with a bank or some other counterparty, you need to understand the implications of that. The counterparty—not the transaction but the counterparty—is typically regulated. You should understand how they're regulated. It's worth going to school on that. It's a much more flexible product. Anything they can model, these guys will sell and they can model all sorts of things that you haven't even imagined yet. However, you have counterparty risk. They may not pay. You should worry about that. These are long-dated instruments. Typically in five to seven years they're volatile instruments. You should recognize that things can go wrong. Pay some attention to the counterparty exposure you would be developing. You also have problems with concentration. You do a lot of this business with a single bank and all of a sudden you have an accumulation problem.

There are also liquidity issues to look at. I would argue liquidity is not as big a deal as you might think. The reason I would argue that is because these are assets you typically want to buy. You don't want to trade them because they're exactly the assets you want to hold to hedge your maturity value.

Liquidity issues arise when your projections are off. In particular, if you are obliged to write contracts with policyholder put options, this becomes more of an issue. With exchange-traded options, exchange-traded instruments, you have a high degree of liquidity generally. Stocks, plain vanilla options, futures, all offer you that particular opportunity. Over the counter is an illiquid market. In many cases the only one who is prepared to buy this asset from you is the person who sold it to you. Your exotic options, swaps, forwards and structured assets all suffer from lower liquidating, and you have to put that into the equation as you consider how to proceed there.

Finally there are RBC issues. The reason why buying stocks and buying a put is a bad idea fundamentally, from my point of view, is the cost of carrying the capital around—never mind the transaction costs and all the things that go with that. It's not a clean transaction; the RBC is high. Puts and calls have a much lower RBC on a much lower number. But if you examine it carefully you probably come to think that it's probably a better deal all around to go that route. Futures and swaps are currently zero, because futures and swaps are off-balance-sheet items. Now, they may not be zero forever. They're not zero in Canada.

Those are considerations when trying to decide what kind of hedging strategy to put into place. You should worry about counterparty; you should worry about liquidity; you should worry about the RBC implications; you should worry about how easy it is to explain to management and whether it's fundamentally doing the job for you on a relatively economic basis.

Ms. Charlene M. Barnes: I'm with the Union Bank of Switzerland. I work on the equity derivatives desk and I'm very involved in a lot of equity-indexed hedging. We have all sorts of products. There are some as Lilia was saying, with universal life, and some immediate annuities are coming out, too. I also work for the insurance risk management group of UBS. We work a lot in the fixed-income derivative area, trying to take care of all insurance companies' derivative needs.

I'm going to talk about hedging equity-indexed products. The definition of a hedge is to protect oneself from losing by a counterbalancing transaction—that's out of Webster. We all know what losing is. Losing is not making any money, or worse yet actually making negative money here. With equity products, you have general product risk. You have the risk of a lapse, you have the risk of interest rate changes. Anything that can happen to a regular annuity can also happen to an EIA. You also have index risks. This is the risk that basically your index will increase with an equity-indexed product and you're just not going to have made enough in your assets to cover it. There's also a bit of correlation risk. This is mostly the risk that lapses will increase when the index is decreased. You're going to lose money on whatever hedge you had out there, and that's how it can greatly increase your lapse losses.

The next thing out there is the definition of a derivative. This is a general definition: a derivative is a contract where the payment amount is derived from something else. For example, EIA payments are derived from the level of the index at maturity or points in between, depending on the design. For us to hedge this liability, you must near it in the assets that you're purchasing.

I work for the Union Bank of Switzerland. I sell equity derivatives. I'm sure you can all figure out where I'm going with this. Many insurance companies are very new to equity derivatives; it's a scary field. Many companies would like to not have equity derivatives.

If you're going to have an equity-indexed product and you want to avoid the D-word on your balance sheet, you really don't have too many options. The one option you have that's a real option is reinsurance. If you can find somebody who's willing to reinsure all the risk associated with the index, you can get away with not having any sort of derivatives on your balance sheet at all. The other way to do this is with synthetic options. This is not really an option. Alan was talking a bit about it. It is possible to purchase a portfolio of the stocks of the index and then buy and sell these stocks according to very complicated formulas based on when the index changes. There are some huge problems with this. There is the RBC on the stocks, there are very high transaction costs. You basically have to be constantly trading in very small amounts of these stocks for this to work. Also, this only works in transactions that are instantaneous. You have to be out there, you have to know what's happening with the index and with volatility and interest rates. You have to be there ready to trade your odd lots of stock.

Another real problem with this is that if the S&P drops like a stone, this just doesn't work. You cannot trade fast enough. Synthetic options were very big with pension plans before October 1987. They weren't so big after that.

I'm going to talk about hedging equity-indexed products using derivatives. We'll talk a bit about basic option pricing, option pricing sensitivity, and dynamically rebalancing versus over-the-counter options. I'm not going to talk much about disintermediation.

The option is the right to buy or sell an underlying security at a predetermined price. This is the call definition we've all heard of. If you have a call struck at 95 on Xerox, which matures in a year, you know that in a year you have the right to buy a share of Xerox for $95.

One of the definitions that's a little more relevant here is the option that's the right to receive or pay the difference between an index level and a predetermined price. The S&P 500 is not a security. It's an index of a security, so you can't really go out there and buy the S&P 500. You can buy all the stocks in the portfolio, but you can't really go out and buy the index.

The strike is the predetermined price that we're talking about. We say a bond option is struck at 100. That means you have the right to buy or sell at $100.

The exercise date is the date on which the option can be exercised. We have a little circular definition there, but that's the date when the owner has the right to go to the writer of the option and say—OK, I want to buy this for this price like it says I can do in the contract. The maturity date is the final exercise date. For example, European options have one exercise date, and that's the same as the maturity date. American options, which can be exercised at any period, have several exercise dates but only one maturity date. That's the last date you can use this option.

If you're long in options, that means you own an option. That means you bought the right to receive the difference between these payments. If you're short in it, that means you sold it, which means that somebody else has the right to come up to you and say—sell me the security or give me the difference between this index and this strike price.

This gets a little confusing. Options talk about premiums. You talk about premium for option. That's the price paid for the option. It gets a little confusing when you are talking about EIAs. You talk about the premium that comes in the door. You talk about the option premium and it gets a little difficult to differentiate it. Generally, I'm going to talk about option premium here. In the money option is one where the price of the underlying is greater than the strike, in the case of call options. I'm basically going to talk about call options here because that's the most practical way to hedge these annuities. It's certainly the one that's used the most.

If the value of the index is above strike price and the option is in the money, you have the right to exercise on that date. You could go to the option and say, "OK, give me this money." An option is out of the money when the price of the underlying is less than the strike. Right now the S&P is about 890. If you had a call at 950, that means you have the right to buy the S&P at 950. You would not exercise that if you had the right to do it. It's out of the money. Out of the money means that the price and the strike are the same.

The forward is the implied price of the underlying on the exercise date. That's a bit of a theoretical concept here. I'm going to get into that a bit later, but basically when you own an option it has an exercise date. The forward is the price that's assumed by the market when you're pricing the options. If it matures in a year, the forward price is the price that the market expects in an arbitrage-free world, the price to be. The sensitivity of the option price is the changes in pricing variables—we have the pricing derivatives of the option prices.

Just to recap, the call is the right to receive the difference between the index and the strike price. Calls are in the money when the index price is higher than the strike price. EIAs have embedded calls. The annuity owner receives the minimum, which is the guaranteed portion that Alan was talking about plus the difference between the index and the minimum. The index minus the minimum or strike price is the call. To hedge this liability you must purchase something that mirrors it.

Who has heard of the Black-Scholes model? Who here has it memorized? The Black-Scholes model, which is a formula that many people have memorized, is used to price the European options. These are the options that can only be exercised on one day. There's a whole lot more behind Black-Scholes that tells us how to price every sort of option out there, and it's called the risk-free assumption. It's kind of a subtle concept, but basically Black-Scholes tells us what kind of model to use when projecting the assets.

Let's say we're talking about an S&P 500 option here. Certain parameters go into the Black-Scholes model that everybody knows about that determine the price of it. Those parameters are the strike price and the maturity. Those are the two things that are determined by the people writing the contracts. The option seller and the option buyer mutually decide on the strike price and maturity price for an over-the-counter option. Also, another parameter includes the discount rates. These are the famous risk-free rates, and we'll talk a bit more about that later.

The dividend rates, stock price, and volatility are all things about the underlying security or index. When pricing European models on the S&P 500, you need to have a dividend rate for the S&P 500, the spot price, and the volatility.

Black-Scholes assumes that stock prices in the future will be lognormally distributed. It's just a plain distribution of two parameters. It has a mean, a standard deviation, and all that other good statistical stuff. The nice thing about it and the reason why it's used for this in the first place is normal distribution is kind of the way people understand random things to behave, so we use a normal distribution. But a lognormal distribution cannot go below zero. If you buy a stock, you can't lose any more than all your money.

The mean of distribution is a forward. Here we have a strike price, so that if the stock of the S&P 500 or the stock price is above the strike price and it's a call, it's in the money.

When you have a distribution with two parameters, you need to determine what parameters to use. The parameters are determined by the Black-Scholes model and the market. The mean of the distribution is the forward price, and the forward price is equal to this formula: the spot price, 890 for the S&P right now, times R minus D times T, where R is the risk-free rate, D is the dividend yield, and T is the time to maturity. This determines the mean of the distribution of the future stock prices. The volatility used for the distribution is the volatility implied by the market.

There is a very large market for S&P options. It's quite liquid for an exchange-traded option. Volatility will be implied by the market because you have all these prices, you have your Black-Scholes model, you can take the prices, you know what the risk-free rate is, you know about what the dividend yield is. You know what the time for maturity is because it's defined by the contract, so you can back out the volatility. The entire market's idea of volatility is what is reflected in the option price. It's just supply and demand.

The Black-Scholes model is good for pricing nice, simple European options. But as soon as you have anything more complex, which is as soon as you start thinking about anything other than the simple point-to-point model in the EIAs, Black-Scholes doesn't really work. The Monte Carlo simulation is used to price model options with no closed form solution. That basically means any sort of option that you can't write a nice, neat little formula for. For the Monte Carlo simulation you use stochastic projections, but you use the same parameters as Black-Scholes. Often when doing asset allocation work, you think—what kind of return over the risk-free rate do I expect equities to get? You don't do that here. You assume the distribution for stock that's implied by the Black-Scholes model. You're not trying to figure out whether you should invest in stocks at all; you're trying to hedge. To hedge, you have to use the instruments available in the marketplace. Those instruments assume the distribution implied by Black-Scholes. If you want to find the real hedge cost, you will have to use those assumptions.

There are a couple practical considerations. We talk about the risk-free rates. The risk-free rate in Black-Scholes is the rate that you use for the distribution. Generally, the risk-free rate used to price options is equal to the swap rate. For example, the seven-year swap rate is the seven-year T-bill rate plus 30 basis points. It's used to determine the forward index and for discounting the values. When using Monte Carlo simulations, you come up with many values and discount them at this rate.

When many people hear risk-free rates, they think—of course, I'm going to use Treasury. But the indemnity and subtlety of Black-Scholes, the entire idea is that you need to use a rate that a high-quality lender can both borrow and lend at. Nobody can borrow at the Treasury rate except for the U.S. government, and the U.S. government doesn't write many options.

The same is with volatility. You need to use the implied volatility of the market. You can't use the historical volatility. It has to be determined by supply and demand. If you use historical volatility which is much lower than the implied volatility right now, you will come up with an unrealistically low hedge cost. When you go to the market and try to hedge that volatility, you won't be able to buy it as cheaply as you think you can.

From the Floor: How does the repo rate compare with the swap rate?

Ms. Barnes: Generally the repo rate is a much shorter term rate, but it is much the same thing. For practical purposes, if you're going to get into serious hedging, especially over short-term periods, you shouldn't count on making money off the difference.

Basically the reason why we use the swap rate is because the big market makers out there (banks and anybody else who is issuing any sort of option), are going to use this Black-Scholes rather than Monte Carlo type simulation. They're going to use the risk that they can really borrow and lend at. The reason why we use the swap rate or the London Interbank Offered Rate curve is because mostly banks are into it, they're very high quality. A lower-quality institution will not be able to play because the option prices that it comes up with will be higher. Your borrow-and-lend rate will be higher. Therefore, interest rates will be higher and will increase your option prices. It just doesn't really work. The swap rate is basically the lowest rate that anybody who is in the financial market can borrow and lend at.

One question I was asking myself when I was writing all this is—why am I telling you all of this? It's a lot of information, it's quite technical, and I haven't given you enough to go out there and price your own options. The reason I'm telling you this is because I'm trying to get across an idea of option price sensitivity. When the parameters of the option price—the strike price, maturity, the risk-free rate, the implied dividend rate, the spot price, and volatility—move, your option price moves. You have two options here. You have the option that you bought and the option that you sold and your liability. The option price determines the participation rate that can be offered. We talked about that. Basically you have a little budget, you pay for the guarantee, you have some extra money left over. If the option prices are sky high, you won't be able to offer that great of a participation rate. Also, the sensitivity of the option price is the same as the sensitivity to liability. That's very important when doing any sort of hedging.

We'll talk a bit more about the sensitivity and we'll talk about the greeks (measures of financial sensitivity). There are a couple different names for some of these things. I'm just going to talk about the four major ones. The percentage change in the option price for the percentage change in one of the underlying variables is the English way of saying a partial derivative. The first one is the delta. The delta is change in the option price, percentage change in the S&P index or whichever index you happen to be using.

The rho is the percentage change in the option for a change in interest rates. The kappa is the change in the option price for the change in implied volatility. The gamma is the change in the delta for changes in the S&P 500 index. Basically it's kind of a convexity of the S&P 500 index.

I will tell you how these work for call options. If you own a call and it has a certain strike price and the index increases, the value of that call increases, so the delta here is positive. When you own a call and interest rates increase, the value of that option will increase. When you own a call and the implied volatility in the marketplace increases, the value of the option you own increases. The gamma is a bit more complicated, but basically when the S&P keeps on rising and rising and rising, and you have a call—say you bought a call that struck at 200 several years ago—every dollar change in the S&P will basically be a dollar change in your option price. So far in the money, they're almost the same thing.

We have these partial derivatives. How do we determine these greeks? If you have a formula, it's really simple. You just get your local mathematic type to figure out the partial derivatives of your option formula. The only real option formula out there is the Black-Scholes which is for European options only. If you don't have something with a nice closed-form solution, you use the Monte Carlo pricing with bumps. You bump volatility up, you bump the volatility down. You bump the interest rate up, you drop the interest rate down. You try to see what the sensitivity is for large changes in the underlying index.

Now we know how these options are priced. We know what the sensitivities are. What are we going to do? You can hedge with the greeks. The greeks are a measure of the change in price due to the change in one of the parameters. You're hedged when the assets and the liabilities have the same greeks. A change in the liabilities will be the same as a change in the assets for any of the changes that underlie. Of course, in most cases when you have very disparate liabilities and assets, you're hedged only for that instant.

This should sound familiar to you. It's the same thing with the interest rate. The interest rate is the derivative of a price of a bond or some sort of fixed-income liability. The change in the interest rate is the same for both the assets and liabilities when your duration matches, so basically you can kind of think of this as duration matching to the fourth power.

I'm going to talk about two hedging methods. This first is replication. This is the one that’s used far and away by most insurance companies to hedge their risk—certainly in the initial stages of the product. This is where an insurance company will purchase the matching assets, for example, S&P options, to offset the embedded optionality in the annuity product. Basically if you have a product that has a seven-year S&P call struck at 110.69, you buy the exact same option.

The other way is by dynamically balancing. This is what we do. This is how we take care of our hedges. When we buy annual discrete lookback options for seven years with a strike of 110.69, we don't go out and find a counterparty who just happens to want the exact opposite transaction. We will dynamically rebalance this. An insurance company could do this. The company would purchase a series of options and futures to offset the embedded optionality in your product.

The replicating transactions mirror the liability exactly. These are over-the-counter transactions. These are contracts you get into with individual counterparties. You set the contract up. You can come up with an asset that has the same maturity, the same index, the same strike as your liability. It just matches it. All these things are the same; they have the same greeks.

If you're hedging and you have $100,000 worth of equity-indexed premium to hedge, how would you do it if you were dynamically balancing? You determine the Greeks of the EIA. You determine the greeks of all the available instruments in the marketplace that you can use to hedge. You would purchase the appropriate instruments and then monitor the position and make the appropriate changes. This is where it's very important to remember that you have to use the Black-Scholes assumptions when coming up with your pricing. That's how the available instruments are priced.

Let's go through a relatively straightforward example. We have XYZ with its EIA product. It is termed at seven years. The minimum return is a standard nonforfeiture. That turns out to be 110.69. The index return is 100% of the S&P 500, which is really good for right now. The surrender provision is no S&P vesting, which means that the index portion is not vested and there is the 8% surrender charge.

How would it actually hedge this if it is replicating? The company would just purchase an over-the-counter seven-year European option with a strike of 110.69. If it is doing dynamic balancing, it would immediately take the premium and purchase some index future contracts and some one-year maturity options, too. These are both things that are available on the exchange. The company can monitor them.

We're talking about straight European options. There are some other ones out there, too. There are look-backs, there are Russians, and there are Asians. Most of these are still on the S&P, so basically it's just a difference in the formulation. The index, instead of just being whatever the number is at the end of the period, is a function of the past along the entire period. Most of these must be priced by using Monte Carlo techniques. With a few of them you can kind of use little adjustments to some sort of other option function, but pretty much you have to use Monte Carlo.

For the replicating transactions you can just go to a counterparty and get the exact opposite or the exact same, I guess, transaction. For the dynamically balancing, you just have to keep track of the different function. You have to find the Greek for a different type of payoff.

A good example is with the look-back-type options. If the S&P were to increase dramatically during the second year, you've locked in that second-year amount, so your Greeks will change substantially at that point because your sensitivity to the S&P has changed dramatically. Why? You're still going to owe money if the S&P suddenly goes way down. That's just the type of thing that you can keep track of. Basically, it's just the same Monte Carlo simulation with a function in there that keeps track of the entire path. A lot more computer time is used.

After having preached the gospel of Black-Scholes' assumptions on your option price, I want to say a word about inexpensive options here. Option prices obviously determine the product design. There are certain designs that people love to do but the options are just way too expensive or the option price determines what participation you can give. It's very tempting to embed a cheap option and get better product characteristics. The problem with going out there and saying, oh well, for 110% participation they'll just sell this thing like crazy—will be really obvious when you're buying over-the-counter options. You're just not going to have enough money. The option price will come at a certain amount. It's going to take up more than you have left over in your budget after you find out the guarantee.

With dynamical rebalancing, there's the argument that option prices are too high, that the implied volatility right now is too high or that the implied dividend rate is too low. There are ways that an insurance company or anyone can hedge this. It's not a perfect hedge, but it tends to work in most situations. At least volatility is low. You can hedge by only using futures and not using the options and just using the formulas to change your position.

The insurance company capital is at risk whenever you do this. Basically you're making a bet that something is wrong with the market. If you win your bet, something was wrong with the market, your hedge scheme worked and cost less money, and you give all your money to the policyholder. If the hedge doesn't work, you're out the money and you still owe the policyholder. Take it very seriously. It's not something that's easy to do. It's something that is very risky. Many equity derivative desks quietly go under every year. It happens.

Just a word on that counterparty risk: you can measure the credit exposure on an over-the-counter derivative transaction. This is very important. I said that a couple equity derivative desks go quietly under. It does not mean the bank backing them up goes under, it just means the bank decided it wasn’t making money and decided to stop. But you do need to keep track of this. For several companies this has become a very significant line of business. They all of a sudden own tons of over-the-counter options. You need to keep track of your counterparty risk. You just need to keep track of the market-to-market value of the options and the value at risk. You can just basically use that compared with traditional credit guidelines. You'll have a mark to market; you'll have an idea of the sensitivity so you can just basically go through the typical credit analysis. There are more factors.

I just want to sum up and talk about the relative risk and expertise of these different hedging methods. The easiest way to hedge an EIA and never lose money off of it is to not write one. I kind of rank these in the order of risk. Reinsurance is out there. Reinsurance runs the entire gamut of possibilities. You can reinsure all your S&P risks theoretically so you have no risk at all to pass onto the reinsurer. Or you can reinsure a portion of it, or you can only reinsure the lapse risk on the interest-sensitive part. It's between you and the reinsurer.

Then there's replication. The advantage of replication, which is when you buy the over-the-counter option that matches, is that you don't have to keep careful control over the greeks. They will basically match at first. You, of course, have to know what you're doing. You have to understand the sensitivity of these things, but your capital isn't really at risk if the balancing behind the over-the-counter transaction doesn't work. You aren't buying the rebalancing skills of the counterpart—you're buying a guarantee to pay them this much at a certain date.

Finally, there is dynamic rebalancing, which is basically bringing risk in-house. Some are doing a blends of a method. Some are basically doing dynamics rebalancing. One in particular is doing dynamic rebalancing but is buying long-dated over-the-counter options. You can't buy a five-year option on an exchange, so it goes to the counterparty and gets a five-year call, but its liability is more complicated than that.