RECORD, Volume 29, No. 1*

Washington, D.C., Spring Meeting
May 29–30, 2003

Session 18PD
So, The Equity Markets Don't Always Go Up? Capital Markets Hedging of Variable Annuities and Equity-Indexed Annuities

Track: Investment

Moderator: MARSHALL C. GREENBAUM
Panelists: D. KENT FREEMAN†
DANIEL R. PATTERSON
JUN ZHUO‡

Summary: This session examines the nuts and bolts of implementing a capital markets hedge program for offsetting the market risk embedded in variable annuity and equity-indexed annuity products. The following hedging considerations are explored:

- *Dynamic versus static hedging;*
- *Full protection versus "tail" protection; and*
- *Delta versus multiple "Greek" hedging.*

The participants develop a better understanding of the infrastructure and expertise required to implement an effective equity hedge program.

MR. MARSHALL C. GREENBAUM: This session is entitled, "So, The Equity Markets Don't Always Go Up?" We focus here on capital markets hedging programs for both variable annuities (VAs) and equity-indexed annuities (EIAs). My name is Marshall Greenbaum. I'm going to be the moderator of the session. I'm also an actuary with a firm based in New York City called Constellation Financial Management Company. For those of you not familiar with Constellation Financial Management Company, in

---

†Mr. D. Kent Freeman, not a member of the sponsoring organizations, is managing director at Constellation Financial Management Company in New York, N.Y.
‡Mr. Jun Zhuo, not a member of the sponsoring organizations, is a principal quantitative analyst at AIG in New York, N.Y.

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

a nutshell, we purchase deferred-acquisition cost (DAC). We purchase DAC, do capital markets transactions where we provide cash up-front for the receipt of a future variable cash flow over a defined period of time, and we take all the capital markets risk associated with those variable fees. Transactions might be paying six percent or something like that up-front for the payment of *x* basis points, *x* basis points mortality and expense (M&E) fee, for a period of, say, seven or eight years.

From when we started in 1996, the company's business philosophy has been to actively hedge out all capital market risk that we deemed cost-effective to do so. At this point, we have a seven-year track record of managing and monitoring what we call basis risk, and that's going to be talked about to some extent by our panelists.

Actually I had a request. Just before we get started, we're curious to know how many folks in the audience attended the all-day seminar yesterday on hedging VAs without tiers. Could we have just a show of hands?

I am joined by three extremely qualified presenters. A colleague of mine, Kent Freeman, also with Constellation, is going to talk about the steps required from a modeling and analytic standpoint to put yourself in a position to utilize a hedge program and the assets, futures and options that would offset your liability positions. After that, Jun Zhuo is going to extend the concepts that Kent talked about by adding a case study going about a static optimization hedge example that he's put together. Lastly, we're going to switch gears a little bit, and Dan Patterson is going to talk about his dynamic hedging experience with EIAs.

Before turning it over to the panelists, I just wanted to make a couple of brief remarks and set the stage a little bit. When you look into implementing a capital markets hedge program, there are two main or high-level options available. One is going out and purchasing a financial instrument that would protect you in the case of market declines, and I lump all these into the same kind of category. You could go out and buy an option from Wall Street. You can enter into a futures contract. You could go out and get guaranteed minimum death benefit (GMDB) reinsurance. When you think about what GMDB reinsurance is, it's more or less a financial guarantee, a put option with some actuarial wrinkles to it.

You also have the ability to go out to Wall Street and design some very structured put options that would more mirror your liability characteristics. On the M&E side, they hedge your revenue risk, have the ability to enter into total return swaps where you exchange your variable fees for fixed fees, and can either exchange your actual fund returns or index returns for fixed known fees over time.

There's another category that I think often gets overlooked in terms of a risk-management strategy—the actual sale of a risky asset associated with the products, and that being future revenue fees over time. If you think about future M&E fees that you're going to collect, it's more or less an asset, whether it's capitalized on your balance sheet or not, and if you can sell that for a known quantity and eliminate the riskiness of that asset going forward, you've certainly entered into a risk-management solution.

On the first category I want to stress the point about going out and getting any kind of derivative or financial product and just purchasing it. Because of the way those derivatives are priced, if you're looking at what it's going to do to your product return in kind of a static scenario when the markets go up, say, eight or 10%, you're going to be in a situation where every time you look at the purchase of a derivatives instrument you're going to be reducing earnings. So if you're in the mind-set that you could utilize a capital markets program, that's not going to have that effect, and you're in a situation where you can't reduce earnings, you can save yourself a lot of time and aggravation because essentially there's no free lunch. You have to keep in mind that you get something for that derivative instrument. You get the benefit of risk transfer.

The point I'm trying to make with the second concept, which is if you actually sell your fees and you sell them in what I call a capital-efficient structure, meaning that when you receive the cash it actually creates capital, you could be in a situation where not only you create risk transfer, but you actually don't do anything to your product returns or can actually increase them.

In looking at any types of capital markets programs, it's always critical to think about all the objectives you're trying to solve. It's a little bit of a juggling act. You want to increase your product returns, of course, from a statutory and a GAAP basis. Of course you want to reduce the riskiness of the product. In today's environment, a DAC asset is not looked upon favorably by financial analysts, and eliminating it to some extent would be viewed as a very positive thing to do, as of course, anything that would eliminate capital today where it's difficult to raise capital at attractive terms.

Lastly, I think that we've been in a situation where many insurers have gotten to a point where they've looked at capital market solutions, but they haven't necessarily been able to figure out the right accounting to go along with it. If they hadn't gotten stopped at the first obstacle, they got stopped by the second DAC obstacle, and what's desirable from an accounting perspective is that you have accounting symmetry. Either mark to market your asset and your liabilities, your liabilities and the hedges that you put on, or you don't mark to market them. If you have one going against the other, you're going to introduce income statement volatility, which has been an undesirable thing to do from an insurer's standpoint. That's under the scrutiny of earnings volatility.

Keeping all those objectives in mind, I wanted to point out that there is a transaction that solves a lot of those objectives. It's something that really hasn't been done today in the VA marketplace, and it's actually to go ahead and sell 12b-1 fees embedded in the product. Many variable annuities today have 12b-1 fees in them, and there's an important distinction between why you can sell 12b-1 fees and you can't sell M&E fees and get desirable accounting and capital treatment. That's because a 12b-1 fee is a distribution fee that's earned at the point of sale. There's no continuing involvement in the generation of that cash flow, and so you could sell those fees that are going to be earned over time for up-front cash that you could book as capital and also take down the DAC on your balance sheet. Because you could do it in such a manner, at such an attractive rate of return, compared to an equity like hurdle rate return, you don't have that economic give-up that you have typically when you purchase outright derivatives.

Lastly, I want to stress the point that when you're looking at these capital markets solutions, it's very critical that you look at all of these competing objectives, and that you also look at them in what I call a risk-adjusted framework. Otherwise, you're almost always going to come to the wrong conclusion, which is basically not to hedge. I think that's what's happened a lot in the past, that they weren't necessarily looked at in terms of risk-adjusted return on capital and how to price this VA. Then there's been a situation where VA writers haven't hedged. Now they're looking at this with the recent market declines and realizing that that's not a very attractive business model on a going-forward basis.

Having said all that, I'd like to turn it over to the panelists. First up, Kent Freeman is going to be speaking. Kent is a managing director and head derivatives trader at Constellation. He's essentially responsible for developing and executing all hedging strategies at Constellation. Prior to joining Constellation, Kent was head of structured products group at Paribas Capital Markets in New York, where he was responsible for structuring, modeling and trading exotic fixed-income and cross-market derivative securities. Early experience includes the equity derivatives and emerging markets trading desk at Citibank. Kent holds a doctorate in applied mathematics from Harvard and received a Bachelor of Science degree in electrical engineering from the University of Cincinnati. Let's welcome Kent.

DR. D. KENT FREEMAN: Good afternoon. The advertised purpose of my presentation is to outline the steps required in order to achieve an effective hedging program for VA products. That's rather a large task to accomplish in 20 minutes, so let me just give you some idea of what to expect. It's not going to be a detailed manual of how to do it but, rather, just a high-level overview of some of the issues you may want to consider when you're migrating from an actuarial-type model to a capital markets hedging model. Second, we're going to take some of those ideas we develop and attempt to apply them to an illustrative example.

These are the seven steps, and we're going to devote one full slide to each step, but let's just go through them very briefly here. Step 1 is computing the net present value (NPV) of the cash flows. Step 2, model the cash flows. Step 3, model the fund returns. Step 4, model the tradable indices. Step 5, compute the sensitivities. Step 6, decide on the target risk profile. Step 7, execute your portfolio of trades. Steps 1 and 2 are essentially the same steps you would do in any actuarial pricing model, and there'll be very little difference in what we do for a hedging model and what we do for an actuarial model. Steps 3 and 4, on the other hand, are your capital markets pricing steps, and you'll find that you have to actually use what might be considered a lot more detail in order to accomplish an effective hedging program than what you would in an actuarial model. Steps 5 to 7 are just the actual hedge calculations and trading decisions you make.

Before I go on, Steps 1 through 4 are in a backwards order. They're in a top-down rather than a bottom-up order. You might think it more natural that in terms of dependence, you first need your tradable indices to calculate your fund returns. Then you need your fund returns to calculate your cash flows. Finally, you need your cash flows to NPV your cash flows. Don't get confused by the dependence. That's just the way I wrote it out.

Step 1 is to compute the NPV cash flows, and what we're concerned with for the moment is assume that we know the cash flows, and we're only concerned with the actual NPV step. That's because of the approach I just described. Anybody can do that. That's very easy. There's only one point we want to make here, and that is: Do you have to do your discounting based on a market-based yield curve, like a London interbank offered rate (LIBOR) that is actually stripped and computed from instruments that actually trade in the market, like deposit rates, Euro dollar futures, and swap rates, rather than some actuarial discounting rate?

Why do you want to make sure that your discounting is tied to the market? When you're hedging, that necessitates both taking in cash and paying out cash, and you have to make sure that that cash is funded at current market rates, not just for the hedger or the insurer but also for the hedging counterparty who's actually pricing the hedge trades. To make sure that your VA cash flows are priced consistently with the hedge trades, you have to make sure you're discounting using a market-based yield curve.

Step 2 is actually modeling the cash flows, assuming you know the mutual fund returns. Again, that's something you have to do in any actuarial pricing model, and it's essentially nothing different for a hedging model. The main task here is to take the policy level data and aggregate it up into cell-level data using representative dimensions. These include market-type dimensions like portfolio allocation, the amount a particular option is in or out of the money, and actuarial-type dimensions like age and sex and so forth. Just choosing how many cells to use is really a key step, because you have to make sure that if you have products that have drastically different market risk behavior, that they get their own cell, and that they're not aggregated together into the same cell because you could actually lose a very important market behavior by doing that. On the other end, you would like to be able to model every policy separately, but for computational reasons that's usually not an effective way to go about it.

Finally, you can include in this framework any type of cash flow that is dependent on the market, such as M&E fees, surrender charges, tax relief riders or guaranteed minimum benefit (GMB) obligations that could be either death obligations or living obligations. Not only can you include all these different things, but you probably want to include as many different market-dependent things as you can, because you'll see that there are actually diversifying benefits of looking at a global hedging of everything at once rather than, for instance, trying to pull off the GMB and hedge it separately. You'll benefit from the diversification of looking at everything at one time. I hope we'll see that in the example that we look at.

Now we get to the capital markets modeling. First step is modeling the mutual fund returns, which is a step that we think about quite a bit at Constellation because that's our main line of business. For the moment, again, let's assume we know some tradable instruments, and we're just concerned with actually modeling how the mutual funds behave. Black Scholes option pricing theory unfortunately only applies to assets that are liquidly traded and can be shorted, and mutual funds are typically neither. They only have a price quoted once per day, and you can't short mutual funds, so you have to find a way of getting around that. The way you do it is take a set of instruments that you can trade liquidly and you can short, and you model the mutual fund returns in terms of those instruments.

One particular model that is interesting to look at is the equation I've written down, which is my only equation in this presentation. This says, let's assume that the excess mutual fund total return—and by excess I mean excess over the risk-free rate—is a linear combination of the excess total returns of a set of tradable indices plus a idiosyncratic term that I've broken into two pieces here, the expected part and the volatile part. Why is this model interesting? Essentially this model just comes from the obvious extension of the capital asset-pricing model, and it has the flavor of ordinary linear regression. You might think of this as just the regression equation, where you can regress the historical mutual fund returns against the historical tradable instrument returns and come up with your parameters. It turns out you probably don't want to do that, and I don't want to recommend that. The whole issue of actually deciding which indices to use and how to set the beta parameters and the alpha parameters is a very complicated problem in itself. That aside, it's a pretty good model for modeling mutual fund returns.

Other models are possible. You could assume that your mutual funds are static portfolios and attempt to model every security in them separately, but that's just computationally infeasible usually. You could also go to the other extreme and, for instance, just use a single index like the Standard & Poors (S&P) 500 and project all your mutual funds on the S&P 500. That also is not recommended, because normally that won't adequately capture the covariance structure of all your mutual funds, and it turns out the pricing of these products is strongly dependent on that covariance structure of the mutual funds. This is a model that will capture that covariance structure, or one possible model. There are other things you could also do.

The second step in the capital markets modeling is to model the tradable instruments. This is probably more what was talked about in the workshop yesterday, modeling things like the S&P 500 and other tradable indices. The first part of that process is to choose a probability distribution, and if you're modeling things like equity indices, even diversified fixed-income indices and currencies, treating them all as a jointly lognormal set is a pretty good thing to do and works pretty well for most sorts of derivatives you would price. The key point is that, once you've chosen a probability distribution, it will have some parameters in it, like the mean and the standard deviation. Those parameters in a hedging model have to be calibrated to current market prices, and there are two parts to that. There's calibrating the mean, which is the easy part. The mean return also is equivalently calibrating the forward prices or calibrating the expected returns.

Black Scholes tells us that what we should use for that is the actual current price of offsetting the risk of that particular index or hedging that index. For example, if I'm thinking of the Russell 2000, and I go out into the market now, what determines this cost of hedging? The market determines this cost of hedging, like the futures market, the swap market or the index swap market. At any point in time, there'll be a cost of hedging implicit in the futures prices or available in the swap market. For instance, for the Russell 2000 I might be able to go out today and swap Russell 2000 total return in return for receiving LIBOR plus 10 basis points. If that's the current market, then the LIBOR plus 10 basis points is the expected return I should assume in my hedging model.

That's the easy part. It's getting the mean correct. Second, you have the standard deviation or volatility, and that's more complicated because you have more than one thing to get correct here. Typically you have to match to a whole selection of European option prices in the market, across several different maturities and several different strikes. You match to the maturities and strikes that are going to be relevant to your hedging program. You also have to have a model that matches reset option prices if you have a ratchet in your GMB structure. It just turns out that if your model matches European prices, that's not good enough to ensure that it's going to match reset option prices, and if you're trying to price ratchets, it's very important to match the reset option prices as well.

How would you go about doing that? Well, there are a whole lot of volatility models out there, including a local volatility service and a stochastic volatility model. I think some of those were probably talked about in the workshop yesterday as well, and that's another topic of conversation. Let me just say that this step is really the foundation of the first three steps, and in the end your model and your hedging program are only going to be as good as your market calibrations. If your market calibrations are off, then your hedging program is not going to be effective.

Now you have a model that takes this market information as an input and spits out an NPV as an output, and before you can hedge you have to quantify your risks. The easiest way to do that is to just take each of those pieces of market information which you've just calibrated, perturb them a little bit, and see what effect that perturbation has on the mark-to-market value, and that will allow you to measure these various sensitivities. The simplest one, which I'm sure you're familiar with, is delta, which comes from perturbing the individual index levels of each of my tradable instruments. I can also perturb my volatilities. That's called vega. I can perturb my correlations. I can perturb each individual yield curve instrument to get my rho, or I can take pair-wise perturbations. I might be interested in what happens to my delta as volatility moves or what happens to my volatility as the index moves, and those actually can be quite important in exotic derivatives.

I'll talk more about this as we go along. For exotic derivatives, delta hedging alone is usually not sufficient, and I certainly would classify VA-type products as exotic derivatives. You may be able to actually get away with not hedging some of these high-order risks and ignore them, but you should really determine that they're small or unimportant in a value-at-risk framework and really have some rigorous justification for deciding not to hedge it. You should not just ignore them and assume they're not there in a hedging program, because they can swamp your delta hedging and become your primary source of risk if you're not careful.

After you've computed your risk sensitivities, you can decide which ones you want to keep, which ones you want to get rid of, or which ones you want to modify. There are two extremes there. One extreme is to do nothing. We've seen that that sometimes can cause problems. The other extreme is a complete risk transfer, which is what Marshall talked about, where you would just sell all your cash flows in exchange for cash up-front or cash over time. Also, the derivatives market allows you a lot of flexibility, and you can basically have a continuum of alternatives between here, and decide what risk profile best fits your company. You can choose a set of hedge trades that gives you any resulting risk profile. There's more than just risk tolerance. Risk tolerance is important, but also important are which types of risks you want to retain and which types of risks you want to get rid of.

Another factor you might want to consider when choosing your target risk profile is the level of basis risk you want to retain. Marshall also spoke about this. If you do an S&P-only trade, even if it's a very exotic, perfect S&P trade, you're still left with the basis risk of your mutual funds versus the S&P. To bring it to the next level in trade, instead of versus the S&P, several different indices better match your mutual funds, and you will get better tracking, but that trade would be less liquid. You can also bring it all the way to the extreme, where you actually do a trade based on the mutual funds, so you have no basis risk, and you would have zero tracking and presumably the least liquidity. These types of basis risk trades are Constellation's area of expertise.

I also want to make the point what you actually decide to trade—the S&P, the basket or the mutual funds themselves—is independent of the modeling stuff we just talked about. You still need to really model it in detail to get that correct covariance structure. If you do your model just using the S&P, you won't get the correct covariance structure, and you won't get the correct deltas, regardless of which thing you actually decide to trade in the end. The modeling decision and the trading decision are two separate things. The final sort of risk consideration is whether you want to be completely hedged from the outset or you just want to hedge your tail, protect yourself for any catastrophic market event, and we'll talk about that a little bit later, too.

Finally we get to actually executing your trades once you know your risk. I've listed or broken it down into three alternatives in Chart 1. First is delta hedging, which we've said means just trading in the underlying instruments, no options, and the main limitation here is the Black Scholes theory. This says that you can do that, but you only lock in what you think you're locking in, your NPV, if all your other model assumptions that you don't hedge turn out to be exactly right and realized over time, particularly volatility. If you're just delta hedging, you're sort of betting on your ability to predict the other parameters. Alternative 2 is to do a portfolio of options, either European options and/or reset options. Alternative 3 is to do a structured transaction, and what I mean by structured transaction is something linked either directly to the actual cash flows themselves or to some very good proxy for the actual cash flows, and we'll see an example of that.

What are the pros and cons of each alternative? Well, delta hedging is fairly straightforward and easy to understand, but it's also a lot of work because your deltas are constantly changing over time. It requires you to constantly rebalance, to keep track of all your hedge trades and book them and reconcile them, and to perform a full implementation of all the concepts we've just discussed in order to get the correct delta in the first place. If you leave out any of those steps, if you make shortcuts, you won't get the correct delta; without the correct delta, your delta hedging is not going to be effective.

After all this, you're still left over with your exposure to your volatility assumptions that you haven't hedged out, and also your hedge costs. You haven't locked them in over time, and it's really hard to predict up-front exactly what your hedge costs are going to be over the life of your delta-hedging program. If you go to the second step, an option portfolio, certainly you have better immunization against your volatility assumptions because presumably that's how you chose your option portfolio, to balance out your various volatility assumptions. Also, it probably will stabilize your delta, so there'll be less rebalancing required over time.

On the other hand, it also requires a full implementation of all the ideas we discussed, otherwise how are you going to know what options to trade? You really need to go through that whole volatility modeling process to figure out what are your best options to trade to hedge your particular product. Here I've also said, "or faith." You can also have a consultant or an investment bank come in and tell you which options to trade, but then you're basically relying on their model. Somebody in the end has to do the modeling in order for you to put on a trade like this. Also, these types of trades tend to work very well for certain amounts of time provided the market doesn't move too much, but as time passes and the market moves, you may still be required to perform a delta hedge or to adjust your vega hedge over time. Compared to delta hedging, they're higher bid-offer spreads.

Finally, for a structured trade, the effectiveness is self-evident, again because it's tied to either the actual cash flows or a very good proxy for the actual cash flows. You have a very good idea in which situations you're going to make money and lose money over time. You don't really need to understand all those other high-order Greeks or fully implement a model, provided that it's a well-structured hedge. The con for the structured hedge is it's going to be more complicated to document, more complicated to understand, and it certainly will have a higher bid offer than the option hedge, although it may very well still be worth it when you look at your risk versus reward on a comprehensive basis.

Chart 2 shows the example. I've created five equal-sized cells that are dimensioned by differing GMB characteristics, as well as a differing portfolio allocation. The first of the five cells is a ratchet. The second of the five cells, the GMB, is a five-percent roll-up. The third one is the greater of the ratchet or roll-up. The fourth one is a return of premium. The fifth cell is the same as the third cell, except it has a different portfolio allocation, which is more aggressive than portfolio allocation 1.

Among the other assumptions I've made here, the total assets are $1 billion. That's the scale of the numbers you're going to see. The total account value is $1 billion. Also, all five cells include some common features. They all have a 140-basis point M&E. They all have a standard declining surrender charge. They all have a tax relief rider, a 40% tax relief rider capped at 200%. We haven't laid that out here because those features are common to every cell, but they are included. With regard to the two portfolio allocations, I've used seven different indices, five equity indices and two fixed-income indices, a high yield and a Treasury index. As I said before, portfolio allocation 2 is much more aggressive than portfolio allocation 1, and you can see the consequences of that as we go on.

I've expressed the mark-to-market results in two different formats, by cell as a function of account value, and as the basis-point equivalent of M&Es to express that same mark-to-market value. The M&E fees, the surrender charge and the tax relief rider vary by this final portfolio allocation but otherwise don't vary because they're the same for every cell. The difference in the cells is the different GMB characteristics, and that gives rise to quite a large variation in value there. Also notice that for the tax relief rider and the GMB, just changing the portfolio allocation changes the value quite a bit.

I'm now going to go through a bunch of sensitivity graphs, and I'm going to do it rather quickly and try to make a couple of points about each graph. I want to tell you in advance what you're going to be looking at. I hope it'll make sense. There'll be two versions of each graph. One graph will be the total of all cells by cash-flow type. Essentially it'll be a bar chart, and there'll be one bar for every number in this row. Then the second graph will be the GMB only, and it will show it across cells. There'll be one bar for every number here. That's what we're going to be looking at in the next few graphs.

Charts 3 and 4 are entitled, "First factor profit and loss (P&L) profile," and they simply show the variation in the mark to market with respect to a change in the S&P 500 return. First factor here means we've taken all seven indices, and we've projected them onto one dimension, and then we parameterize that dimension by the S&P 500 return. That lets us visualize the delta and the nonlinearity in one dimension instead of seven. It's easier to get a conceptual feel for it looking at one dimension rather than seven dimensions. What's to note about this graph? The tax relief rider has a different slope than the others. That's because a tax relief rider is essentially the selling of call options, and that gives you a nice diversification benefit with respect to the GMB.

Also note the nonlinearity inherent in the GMB, which will even be more evident in the next graph. This again is the GMB-only by cell, the different types of products, and you can see not only the nonlinearity better, but also the variation in nonlinearity across the different types of products. Of particular interest here is the ratchet that at time zero, at zero percent, virtually has no risk. There will be no delta hedge associated with it, but then you become short if you move in this direction and long if you move in this direction. That has the largest gamma position in the ratchet.

We'll pick it up pretty quickly to go through these last few graphs. Chart 5 shows the delta by index, and the conclusion here is the $1 billion of account value translates into roughly, if you add up all the indices, something like $80 million of index delta required to hedge the $1 billion of account value. Also notice again that the tax relief rider has an offsetting benefit for each of the indices. In fact, the GMB almost exactly offsets the tax relief rider.

Chart 6 again is the GMB by cell. The main point here is the large variation in delta across the different types of cells. Again, that's sort of a good indication that you need to make sure your cells are stratified properly when you're aggregating your policy level data. Again, the conclusion is that the $1 billion in assets translates into 800,000 of vega. How much is 800,000 of vega? Is that good or is that bad, the 800,000 short vega position? That basically means if you delta hedge only, and you misestimate vega by five percent, for each percent you misestimate vega, that'll cost you $800,000.

Just as an example, S&P historical volatility has ranged from seven percent to 25% over the last 20 years or so. You can get significant variation in that, and you have to get it right if you're only going to delta hedge. Again, in Charts 7-12, you see that vega also varies by maturity, just not by index, and that's important to note, because in constructing a vega hedge not only do you have to get the right indices, you have to get the right maturities and the right strikes. Again, that varies according to the different cells.

I wanted to talk a little bit about the structured trade (Chart 13). If you take a large number of scenarios and sort them by percentile, you'll get a graph that looks something like this. This is just for a GMB obligation, not for all the cash flows, to illustrate the concept. By structured transaction, what we would mean is just taking the actual formula that you've developed for the cash flow in terms of the fund returns and assuming that the actuarial assumptions are certain and actually writing an option on the NPV of those synthetic cash flows. You could choose with that option exactly how much risk you would want to take, and you would know exactly again under which scenarios you're going to make money and under which scenarios you're going to lose money.

In conclusion, derivatives modeling and trading expertise are generally required in order to achieve an effective hedging program. Different approaches are possible, depending on the amount of resources you want to allocate, and also on your risk tolerance. Structured transactions may provide a solution that doesn't require the full implementation of a hedging model.

MR. GREENBAUM: Next up, Jun Zhuo is going to continue the discussion on VA hedging. Jun is a risk analytic principal analyst with American International Group (AIG). He's responsible for all aspects, quantitative analysis and implementation of AIG's market risk initiatives. This includes analyzing and developing efficient hedging strategies of financial guarantees embedded in insurance products. Before that, Jun was a global market officer with Chase Manhattan Bank and, prior to that, a research associate with Gifford Fong Associates. Jun has a doctorate in theoretical statistical physics from Boston University and also was a postdoctoral fellow in computational biology at the University of California, Berkeley.

DR. JUN ZHUO: Kent just gave an overview of different capital market hedging strategies that can be used to manage the risk in VAs. My presentation on hedging equity risk is focused on what Kent termed alternative tools, how you select portfolio options that can look at the portion of your risk distribution, select a portion of what you don't like, and then, looking at these treatable options, decide exactly what method you use to select them in order to replicate that risk. The method you're actually going to use obviously depends on what works best for you.

As you know, in AIG with the acquisition of Sun America and recently American General, we have become essentially the biggest VA writers in the world. Therefore, we are very well aware of the equity risk in our VA portfolio, and in our department, which is like a corporate market risk management department, we devoted a lot of energy and time to develop effective hedging strategies to manage the equity risk. Today I'm going to share some of our experience with you.

What is the equity exposure in VAs? As you know, there are different types of guarantees, both death benefit and living benefit. Some are a very new type of living benefit, like the last one, guaranteed minimum withdrawal benefit (GMWB), and there is also fee income that links to equity returns. I'm going to focus on these replication strategies using optimization, give a concrete example of how you're going to do that, discuss what the effect is if you do that and then draw some conclusions.

Let's start with a very quick overview. I think everybody knows this already. For the death benefit you have a roll-up death benefit and ratchet, and some of the newer ones, like earning-enhanced death benefit. If you have gains on your VAs, it will help you to pay the estate tax when the owner dies. Then there are older living benefits like guaranteed minimum income benefits (GMIBs), which upon annuitization will guarantee fixed-income streams. Another type of living benefit that is very popular is a guaranteed minimum account value (GMAV), for example a 10-year return of premium or in 20 years double the premium.

There is something new in the living benefit area. I think Hartford first offered that last summer, and I think that Lincoln National is going to offer that or maybe already offered it very recently. It's called a GMWB. Basically you're allowed to withdraw a certain percentage. Normally it's like seven percent of your purchase payment every year until the total amount of the premium has been withdrawn. It is a type of a complicated put option. The other feature is a step-up feature. For example, every five years you have the option to step up, and if the account really goes up, then you step up your guarantee as well. In addition to that, the M&E fee income is linked to the equity returns. For surrender charges, they are proportionate to the account value. Obviously that's another equity exposure. That really is a brief overview of equity exposure in VAs.

How are you managing the risk? You use different formulas and different criteria. The criteria we used are the following. We're actually most concerned about the tail exposure risk. If we want to spend certain money, we want to target that part first, and I think in most cases insurance companies cannot get their inforce data on a timely basis. It's not very practical to really do dynamic hedging, at least not practical at AIG right now. That's why we try to minimize the need for rebalancing the hedging portfolio. That's why we select this replication strategy using a portfolio of options. We try to use liquid, plain vanilla options so we can minimize the transaction costs, and obviously we have the overall goal to minimize the hedging cost.

I'm going to skip that, but just try to mention that the scenario set we used is a risk-neutral scenario set because we try to have consistent scenarios, both to analyze the capital market traded instrument and the guarantees in the insurance product. If you have a different scenario set, then just start over. I think if you show that to management, you will start with a negative value to begin with. Obviously you need to model policyholder behaviors. The dynamic lapse assumption is extremely important for living benefits. Nobody can exactly tell when he or she will die, but for a living benefit, for example, you have a GMAB benefit. If it's only a few years from maturity, and the market is really down, you know people are going to stick around. Therefore, you really need to model that dynamic behavior, the antiselective dynamic behavior.

The method we used to select this hedging instrument is as follows. First, like any stochastic cash flow, you calculate the present value, all the future cash flows, and you have the distribution. That cash flow can include the fee income as well as the guaranteed cash flows. Then you determine which hedging instrument you're going to use to replicate the portion of a distribution you really don't want. You can select European call-and-put options. If you have a ratchet feature, sometimes you need look-back options to match it better. Then you really need to run an optimization program. For a fixed budget you run some optimization program, and that program will determine exactly how much of each option to buy, really to fit the portion of the target liability the best.

Exactly how do you do that? Charts 14 and 15 have some more details to wake up the actuaries who love math. This is your target distribution you have as your scenario. This is a present value of your liability target on the scenario S, and P_H is price of the capital market hedging instrument like European options. So you have a number of them. At each scenario, obviously, you have a different price today. Then this W_i is how much you are going to buy. The optimizer is going to choose these for you. Obviously you have to set up some objective functions in order to do that. The things that we used I believe are called "regret." Basically, this is the tracking error on the scenario S, and then this indicates the average tracking errors.

Before we do that, you need to define what a worst-case scenario is, basically. That would be the most negative tracking errors. So, we want to minimize these negative regrets, but subject to two constraints. First, a worst case has to be bigger than something. I don't want my worst-case loss to be more than $10 million or something like that. And also a lot of constraint is really the budget constraint, how much money I want to spend. I don't want to spend more than the budget. Under these two constraints, I want to minimize the regression, and then this optimizer is a standard linear programming problem, and then it will select the weight and tell you how much of each option to buy.

That's just the value of the option. This is a very general method. It has been used to derive optimal hedges for GMDB and a GMDB with fees. We're going to show an example of this, too, and then living benefits, even the more recent withdrawal benefit, and it's not only restricted to the equity risk. We actually use this to hedge single premium deferred annuities (SPDAs). There is interest-rate risk and some mortgage cancellation insurance as well, so this has very wide applications.

To give a quick example, I think, for some of those who attended the talk I gave, this is the same example, so I'm going to go through this risk relatively quickly. Here you have a kind of VA issued in Japan. It has a ratchet up to age 75 and then stays there for 10 more years, and then after age 85 it will become no guarantee anymore. Suppose the fee to charge for this benefit is 45 basis points. If we want to hedge this, we use the same method we just described.

Here are some underlying assumptions. In this example we use just simple equity indices, but this method obviously is not constrained by that. You can use as many indices as you want. In this case, the lapse and mortality are deterministic, and then you have some policy data of representative cells based on inforce data. Then obviously the market-related data has to come from capital markets, and the model is a pretty standard capital market model. Now I'm going to talk about that.

To start with, if you look at Charts 16 and 17, this is your distribution for your death benefit. If you run 2,000 scenarios, that's how it looks. You want to spend 10 basis points to hedge it, just like you would for reinsurance. If I spend 10 basis points to hedge this, that's how it looks. You can spend a different amount of money, and then you have a different distribution. If you use a capital market scenario set or a risk-neutral scenario set for both the asset and hedging instrument, obviously they have all the same means, just a different distribution. That's all. Hedging costs are included. They all have the same means and obviously different distribution, depending on which one you want.

Is that effective? If you're only concerned about the death benefit, obviously it's pretty effective. As we all know, this is not the only equity risk. There's also fee income on that. If we include the 45 basis points fee income there, your distribution is moved to the right side, and this again is just the same distribution plus the 45 basis point fee income. Then, if you use the same optimal portfolio but just target the GMDBs, then you will see this distribution. They pretty much look similar. So, it's really not hard if you consider all the equity risk.

I think Kent also mentioned the reason that you need to consider all the market-related cash flows together. Some of them have negative correlations. In this case, for example, we have a ratchet, and then we have a fee income (Charts 18 and 19). So, for ratchet, obviously if you have the market going up and then down, you have the biggest exposures, but it is somehow compensated by the fee income because when the market goes up you also collect more fees as a type of offset. You need to take that into account. Just to show this more clearly, if we look at the worst case, with 50 scenarios, GMDBs or GMIBs, so that's how it looks, and then the fees look like this. Under the worst-case scenario, the fee is not the worst. Some of them actually have pretty good fees.

There's natural offset here. You need to take that into account. If you do that, you still spend the same amount of money. Obviously this is the same distribution. Nothing changes. But if I target this GMDB distribution, you can see that. If I spend

10 basis points, and we spend 20 or 30 basis points, you can pretty much eliminate all negative exposures to equities. You spend the same amount of money, but you have a different target here. Here is the example to give you some idea of how it works. This is actually the method we use to really develop some hedges and really put them on.

There are some practical issues when you do hedging. I don't want to go through them, and they are in the handout anyway. I'll try to quickly compare this methodology. What are the pro and cons for this methodology? Obviously this is a very general method. You can tailor it to hedge the select portion of your exposure. First, you have 5,000 scenarios. I have a distribution. I don't like the worst 200 of them. I can target that part and then replicate that part. Then, since we will use portfolio options like a limited buy and hold, you don't have this operation overload there, and obviously with a total capital market there is very limited credit exposure; unless you go to reinsurance you are also exposed to the credit risk there. The advantage for this obviously is that, if the lapse assumptions are wrong, then there's going to be slippage there, and then there's some basis risk obviously. The earned return and credited return are different. Transaction costs can be higher than dynamic hedging, but the problem with dynamic hedging is you really don't know how much is the true cost for that.

For a quick conclusion, if you only carry GMDB like the reinsurance, it's not very effective because of negative correlations between different links to equity risk that are not all taken into account. If you use the capital market hedging strategy we just described here, sometimes it works really well.

MR. GREENBAUM: Next we have Dan Patterson, who's a fellow of the SOA and is an actuary with Allianz Life. Dan has over 10 years of experience in the insurance industry, and his area of expertise is in product development and risk management of deferred annuities, with an emphasis on EIAs and VAs.

MR. DANIEL R. PATTERSON: Well, one advantage EIAs have over VAs is it's kind of the opposite. In the 1990s, with the stock market flying up, everybody forgot on the VA side to hedge, but on the EIA, hey, you'd better be hedging because you're paying off that cash with almost certainty. When I first did this, I thought, so the markets sometimes go up. Most EIA people are fairly active hedgers, but as I thought about it, there are different kinds of hedge strategies that people can take, and they don't always give the same results, depending on different path outcomes. I thought I'd just spend a little bit of time on a simplified EIA product, discuss a few different hedging strategies, look at their performance as a function of market returns, and, lastly, try to combine an EIA hedging strategy with a possible VA block.

I'm going to cover a very simple annual reset crediting strategy that's very common out there right now. The company gets to reset the cap every year to help control future price uncertainty on the options that they use to hedge. Let's say everybody leaves at 10 years, once you work through something else. If I'm assuming people leave at 10 years, I can assume they leave at nine years, and at eight years. I can develop a laddered strategy depending on my expectation of when people leave. We'll just leave it at everybody leaves at 10 years, and focus on that cell.

This is a simple Asian product, at 100% participation and 10% cap. Then I thought I'd throw this in. The nonforfeiture rules are changing on the EIA, I think punishing EIAs too much. They're going to force an 87.5% floor. Assuming interest rates do start to rise some day, it would be 87.5 growing at three percent annually. That's what I'm assuming my minimum floor is. Chart 20 shows what happens. You get your premium, 100,000. You know that out 10 years everybody's going to leave. Bottom line, I need almost 118,000 in cash just on that minimum growth, and then I have this other question mark that I don't know about out there at year 10, which is going to be a function of what the market does and on top of what the market does, since it's an annual reset ratchet-type product that's going to be a path-dependent function.

There are some strategies I could do to hedge that uncertain question mark out there at year 10. First I could go to the over-the-counter and ask them for a long-dated 10-year cliquet. I'm going to not bother even talking about that one. I think it's too expensive, and, frankly, it doesn't match my option on the company to reset the cap if next year's option prices do vary a little from this year. I'm going to say I'm not going to look at it. Let's consider going naked. Let's consider a concept that I call hedge above the minimum required growth that's embedded in that guarantee, and another concept that we've been using at Allianz that we call hedge the option budget year to year, and then I want to get a little fancy and talk about some things that we're looking at and trying to investigate more and get our management comfortable. For a little play on the words, that's half-dressed at the variable company.

Then for all these I'm comparing, let's look at what the 10-year profit and loss (P&L) is, and what I mean by P&L is, let me just assume that premium I get coming in the door gets me five percent. I'm hoping interest rates are going to rise. I'm getting tired of these 3.25% 10-year rates. It's risk-free, no credit risk. I'm going to assume 10-year capped options costs 4.21% in notional, and let me just for the moment assume there's no future price uncertainty. Roughly that's about 22% volatility and a certain slope to the volatility structure. Now the reason I'm going to go ahead and say there's no future price uncertainty is because I'm going to use my management action to move the cap around in the future if somebody comes to me and says, hey, options now cost 4.40%.

FROM THE FLOOR: Is that slope with respect to time or strike?

MR. PATTERSON: That's strike. What I mean by skew is that from your one percent, you go out of the money, multiply that one—you know, drop it 2.4%, roughly.

FROM THE FLOOR: So you have lower volatility when the strike price is 10% out of the money?

MR. PATTERSON: Yes, about two percent, 2.4%. Then my P&Ls follow. At year 10 I'm going to look at what I have left. What do I have in the bank after I do all this? Basically my premium's going to grow. I'm going to pull some out, buy options with it, and get to year 10. My options are going to do something, pay me off something, then I'm going to pay off the liability, and then I have my profit left there for me.

To set the stage, a naked position is shown in Chart 21. Premium is going to grow. I'm going to take my entire premium and throw it in that five-percent earning asset. It's going to grow 10 years. Then I'm going to have to pay out premium times 87.5, that guarantee portion, plus something else. Then you get the line where you go naked. You hope obviously the market falls every year. You make a ton of money. Historically this has been a very poor strategy. It's failed in all but two of the last 30 years when I just back tested against the S&P. In 1971, people would have made money. In the 1990s, they would have lost money. The point is that the standard deviation of the results exceeds 120% of the average gain in this kind of a strategy.

The next strategy I want to talk about is a simplified hedging strategy that says, "Look, at every year I have this guarantee that I have to deliver." This is illustrated in Chart 22. At issue I have to stand and deliver 118% of premium over 10 years. You can back that in. That just basically says every year I have a cost of 1.67%. I don't have to buy an option in that little area between zero and 1.67%, because I'm self-funding that as I'm going. I'll go out the first year and I'll buy an option, struck out of the money at 167. Next year, the market's going to do something, and what it does is going to change my behavior of my option. Let's say that in the first year the market actually did five percent. Now I've credited five percent to the policyholder, and still have to promise to return them 118. Again, over nine years that's going to equate into an annualized cost of 131. Again, I'd say, okay, the second year I'll enter into a position now that struck out of the money, 131 capped at 10%. Does everybody get kind of the flavor of what I'm doing? As I go, I do a different position, and it's basically, how far out of the money am I starting my strike?

Chart 23 shows the results of, I believe, 1,000 scenarios when I try that strategy. Obviously things got a little better. I definitely have tightened down my volatility in my P&L. My average P&L now is about 8.25% of premium, and the standard deviation is still significant—I would say 20% of average profit. I realize that 8.25 has to pay commission. It's not really a time value number. It's just a number that's sitting out there, but you can simply think of it as a profit number that I have to use for profit and expenses.

Obviously comparing that against the naked, it's a much better position when you compare the two in Chart 24. Okay, this looks good. This is one strategy. Then I thought about just discussing the strategy we've incorporated in some of our product design that I call hedger option budget. When you think about hedging, I always think it doesn't matter where you're at in time. You always ask yourself the following question: How much money do I need to make myself replicate my 10-year cash flow? Can I develop a hedging strategy right now so that, no matter what payoff the stock market takes, I will get enough money to exactly give me that 10-year cash flow?

It's kind of a replicating strategy. One can always at any point in time say, "What is it? What do I need? How do I get there?" Really, it's like a tree problem, and you're always basically saying, "Okay, I'm at time T. At time T plus one I need some hedges, and they have value to them. I have to make sure I can get there with what I have today. If I can get there with what I have today, I can get there next year with what I have next year, and so on, until I get to year 10, and that option position or that hedging position will deliver me the required cash flow."

To do this, define a current hedge. We say, let's look at entertaining my two most material positions, so it's kind of a dynamic hedging problem. We're not trying to hedge all the Greeks, like when people think of a full dynamic hedging problem; we really are looking at this as what kind of one-year options can I enter into that will keep this option budget hedged through time? Because of the product design we have where we can actually adjust caps on a yearly basis, we actually use that option to hedge the other risks that would be present if someone didn't have that option. It is a very intensive computational problem, but one that we feel has good success.

Chart 25 looks at the same 1,000 scenarios under this kind of hedging approach, and you see the 50th percentile of 8.64%. I believe the other was 8.25%. I don't know if that's statistically worthwhile to talk about, but the standard deviation definitely went from 20% to less than one percent of average profits, so, there was a much tighter fit around your average level. Again, if you compare this to the prior one, you obviously can see now a definite improvement in the results.

No one laughed about my play on the words. Did anybody not get that? I mean, you had the naked option. Then you had half-dressed. How can we apply this kind of concept to actually working with our variable company? Maybe we're going to sell three billion of EIAs next year, and then we have maybe one billion of variable. I don't know exact numbers, but we have two sizable blocks of variable and EIA that we can actually start to think about. Is there any use in cross hedging? Does it make any more sense to bring the two together, or is there really nothing to add, and should we just do them as if they're independent companies?

Well, one thing about this hedging your option budget, you can start to entertain this concept, and the whole idea about it is laid out in Chart 26. Everybody knows the stock market does really well, and there's more M&E income on a higher base. Variable companies tend to have this kind of leverage line to the market. You hope to God the market does more than your eight-percent expected return, and you'll win. If we start off like the last three years, we're going to end up somewhere down there. So, that's what the line looks like on a stand-alone basis, and the question I'm asking is: With these excess returns over here north of the 60% return axis, can I somehow incorporate that excess money into this dynamic hedging problem on my equity-indexed side? I have excess profits on the variable. Why not, instead of buying as much call options on the EIA, fund it with this? The money I save by not buying the call options on the EIA will help ramp up my line south of the 60% axis.

When I threw them all together this was sort of the line I got (Chart 27). I said, okay, I'm going to modify my EIA hedging problem to say I don't need as much money out in year 10 as I thought that I originally did when it was a stand-alone company. My option budget actually can drop, because I'm going to be getting excess money from my variable company out in those times. I don't need as much option budget. Therefore, you'll bring down the cost of your options, and then if the market were to fall, I have all this excess money that I did not spend on options, that I put in the bank and it earned five percent. Then it can help fund that variable downside.

This is the way that you bring these two together, and you try to figure out, for every 100 million of EIAs, what's the amount of variable I need to do to try to get a good mapping of them (Chart 28)? Obviously there's going to be basis risk involved in this problem, and that's why the blue dots are scattered more than the green. The green line is following that underlying S&P function on the x-axis. The blue line is following something that generally moves about 60% with it, and there's a regression line there. It's not perfect, but it's one that we are moving forward and saying, all right, let's get an understanding of how it works. Let's go slow. Let's figure out what's all there in the space. Let's start monitoring our funds and trying, like these guys are saying, to find some other indices that can make that blue line tighten down even more. That's all I have.

FROM THE FLOOR: I guess that Kent had talked about projecting funds, essentially using capital asset pricing model (CAPM), and someone this afternoon was projecting them at the swap curve, if I understood both gentlemen correctly. I was wondering if maybe they'd want to debate the two approaches.

DR. FREEMAN: The question was what was the relationship between what I said about projecting fund returns using CAPM and what Jun said about using the swap curve. I think, if I understand correctly what you're referring to, they were two separate things. I was talking about projecting mutual fund returns in terms of indices using a CAPM-type approach, and then I would model the indices in terms of the LIBOR swap curve the same way that Jun did. I think we're talking about two separate things here. Ultimately we're both projecting the tradable quantities using the LIBOR swap curve as a basis.

FROM THE FLOOR: I'm sorry. What were you doing with the CAPM exactly?

DR. FREEMAN: That was just to show the relationship between the mutual funds and the tradable indices.

FROM THE FLOOR: That was trying to get the mapping?

DR. FREEMAN: Yes, the mapping of which indices to use.

FROM THE FLOOR: This question is for Daniel. I was trying to get more information on the hedging option budget, what you're actually buying to back that liability. I didn't quite understand what you were doing.

MR. PATTERSON: The question is, when hedging the option budget problem, what kind of positions do we enter into? They're S&P call options. They'd be struck probably at the money for a certain notional amount, up to a certain strike level, and then we'd buy an additional option struck at another, probably up to the cap. The reason I was using two options is there's some curvature in the line you're trying to hedge from year to year. It's not a pure straight line where you could just say one option does it, and you can hedge it with one option line. If you have enough volume, you can get a second option at decent prices and help tighten down the air.

FROM THE FLOOR: This question is for Jun. In your optimization problem, how would you set your time horizon?

DR. ZHUO: Right now you can project as long as you want, and obviously there's some limitation like what kind of a tradable instrument is available. Fortunately for AIG, we do have something like a structured derivative trading desk that can do long-term equity options, but obviously it's going to be very expensive. For this specific example, we just go to 10 years.

FROM THE FLOOR: I have a question for Marshall. I'd be curious if you guys are familiar with how GAAP is applied to all this stuff currently. If I were to sell you some kind of an M&E call option, and the market is up, I'm going to get excess fees. I want secure ties for the next seven years of fees. I'll sell you an option. I don't have to mark that to market. Is that a correct statement?

MR. GREENBAUM: There are a couple of forms of the transaction that could be done to transfer that risk. The most simplistic form, which is to basically write a derivatives contract, and what I referred to as one of the things that you want to contemplate, is trying to have accounting symmetry. That unfortunately does not provide you the right accounting symmetry because you would essentially enter into that swap or that derivative, which is a financial accounting standard (FAS) 133 instrument, an FAS 133 derivative that would have to be marked to market. Then unfortunately you cannot capitalize your M&E fees on your books as an asset that you can mark to market.

In that particular instance, no, but we are proposing a sort of outright purchase and sale of fees, where essentially we purchase those fees, and you essentially receive the up-front cash amount as an asset or capital. Whether it's done on a funded or an unfunded basis, it would just be the difference of recognizing, say, a note receivable or up fronting cash as an asset. The reason why it works from that perspective is we do a transaction that fits a lot of accounting definitions. This enables us to get what we call true sale accounting, which means that you do not have any associated FAS 133 instruments to worry about. It's completely not a derivative type of contract.

It's essentially what we call a purchase-and-sale contract. There are some complications with it, and that's why I said the best way to go about it is to actually sell what we call 12b-1 fees as opposed to M&E fees, because one of the accounting criteria for a true sale to get the immediate income recognition is that you can't have continuing involvement in that cash-flow stream. An M&E fee qualifies as continuing involvement, because you need to be there every year to perform your services to earn the right to that fee.

That is very different from a 12b-1 fee, which is characterized as being earned at the time sale is made, and it just happens to be paid over time. That's the key distinction. For policies that don't have the 12b-1 fee, there's the ability to incorporate a 12b-1 fee and sell it on a going-forward basis. For some companies that have a 20-25 basis point fee, currently there's the ability to sell it on the inforce and not have any type of FAS 133 issues to worry about.

Chart 1

Hedging Strategies: Pros & Cons

	Pros	Cons
Alternative 1: Delta Hedging	■ Conceptually and Operationally Straight Forward	■ Requires Full Implementation ■ Requires Constant Rebalancing ■ Exposure to Volatility Assumptions ■ Unpredictable Hedge Costs
Alternative 2: Option Portfolio	■ Better Immunization to Volatility Assumptions ■ Less Rebalancing Required	■ Requires Full Implementation (or Faith) ■ Delta Hedging may still be Required ■ Higher Bid-Offer Spreads
Alternative 3: Structured Trade	■ Effectiveness Self Evident ■ Full Implementation may not be Required	■ More Complicated to Document & Understand

11

Chart 2

Example: Cell Definitions

Cell	Description	Abbreviation
1	Ratchet Portfolio Allocation 1	Ratchet
2	Roll-Up Portfolio Allocation 1	Roll-Up
3	Greater of Ratchet/Roll-Up Portfolio Allocation 1	Max PA1
4	Return of Premium Portfolio Allocation 1	Ret Prem
5	Greater of Ratchet/Roll-Up Portfolio Allocation 2 (More Aggressive)	Max PA2

12

Chart 3

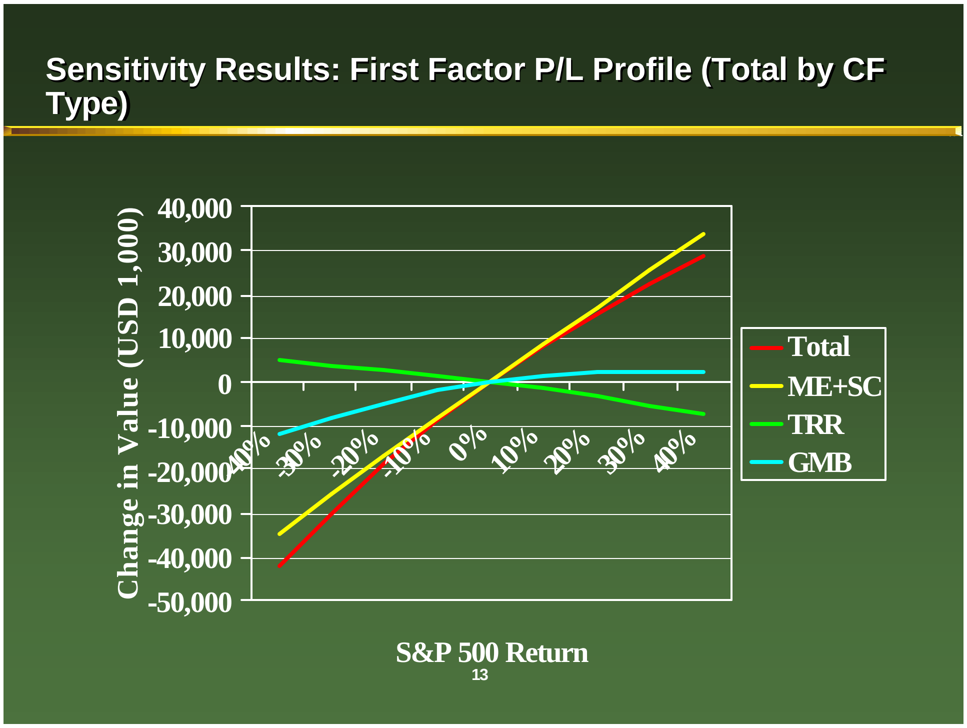

Chart 4

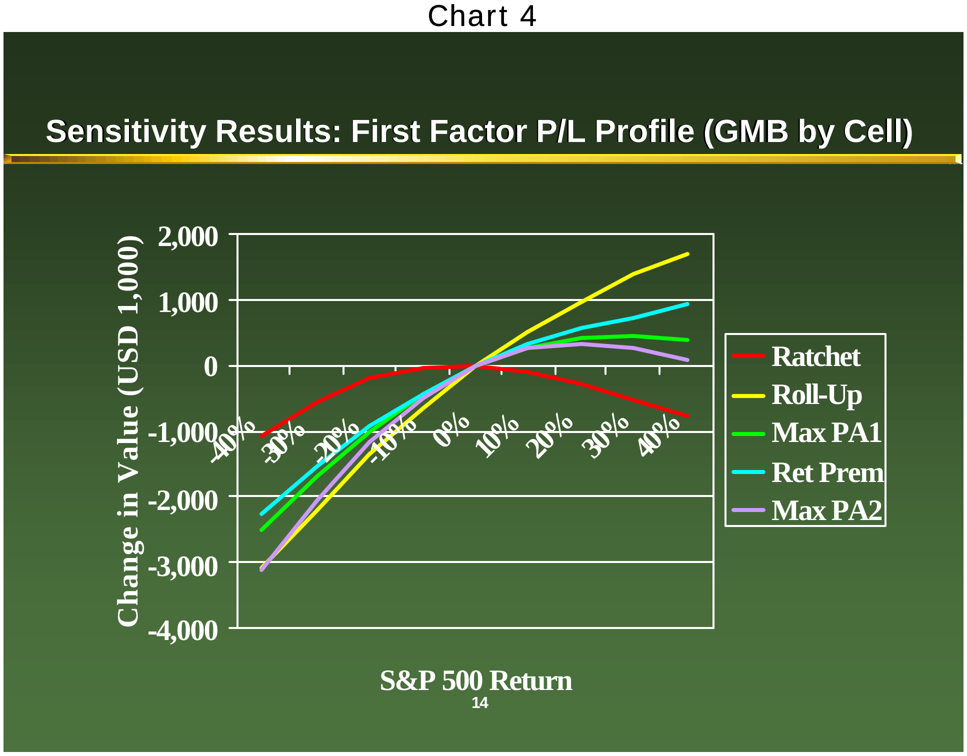

Chart 5

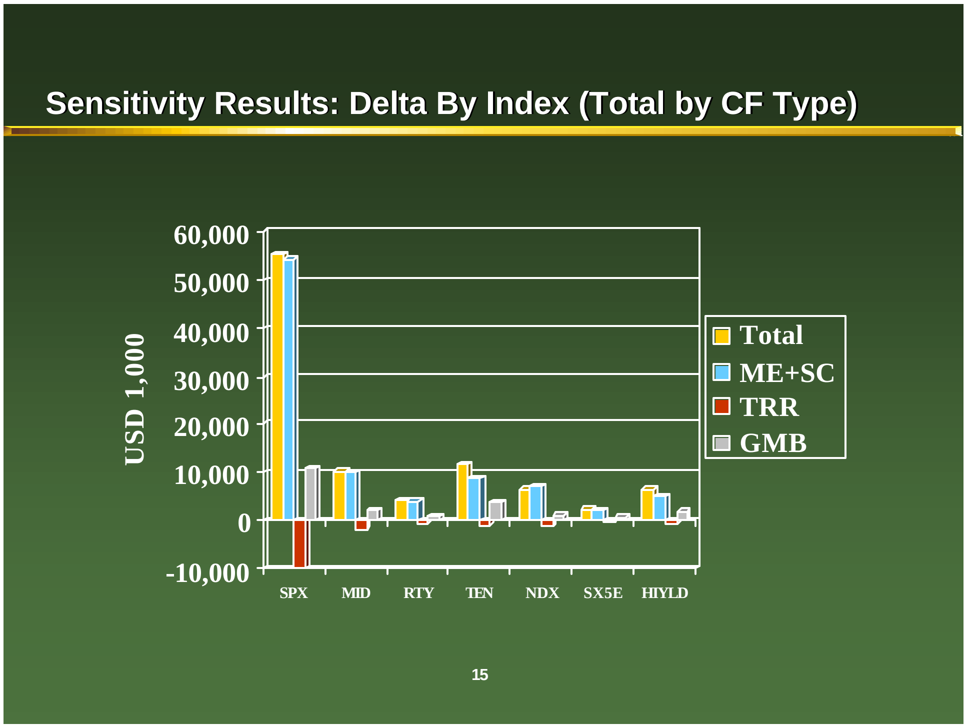

Chart 6

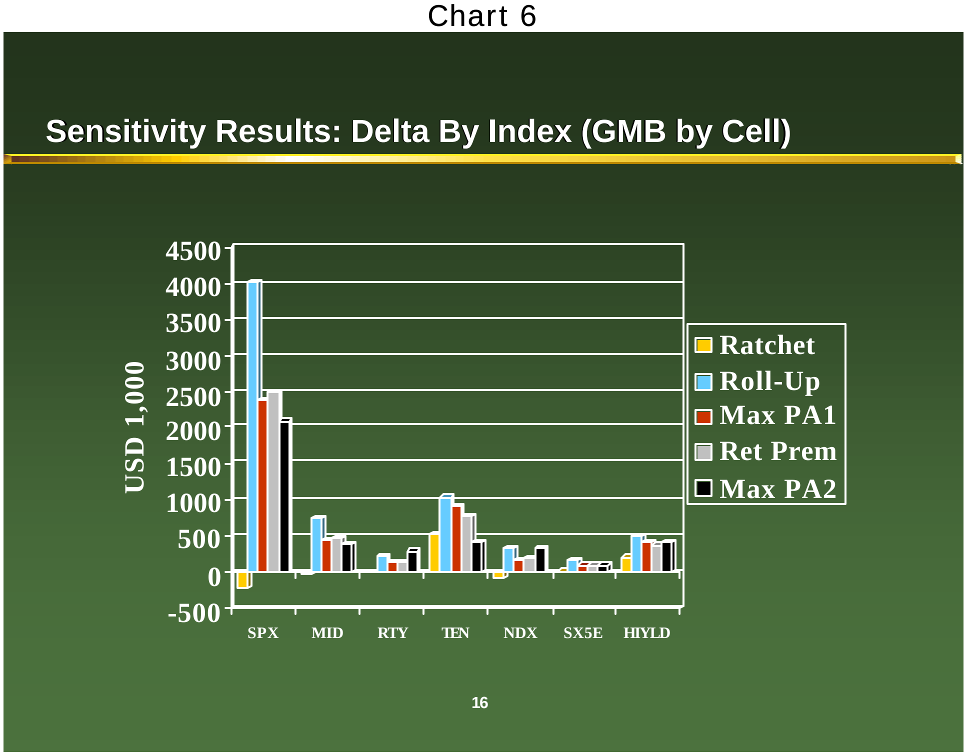

Chart 7

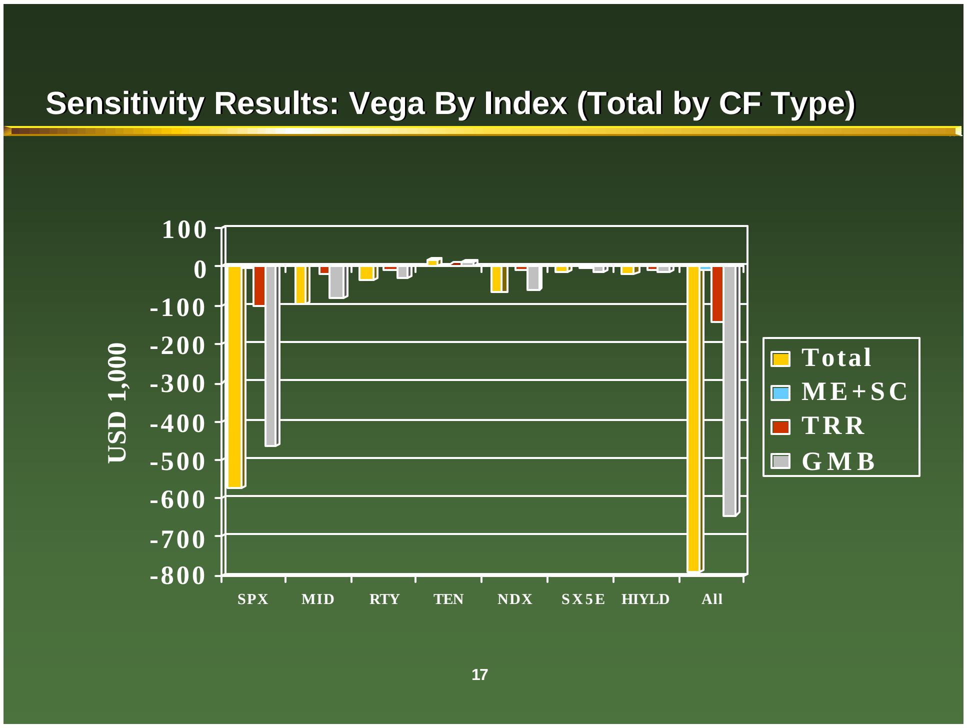

Chart 8

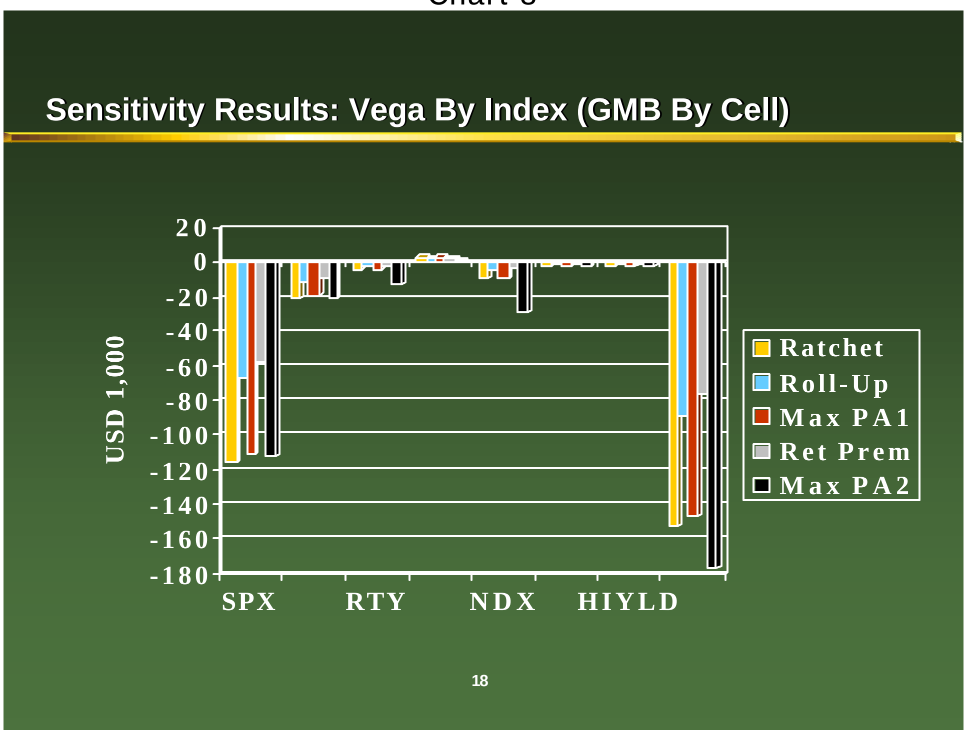

Chart 9

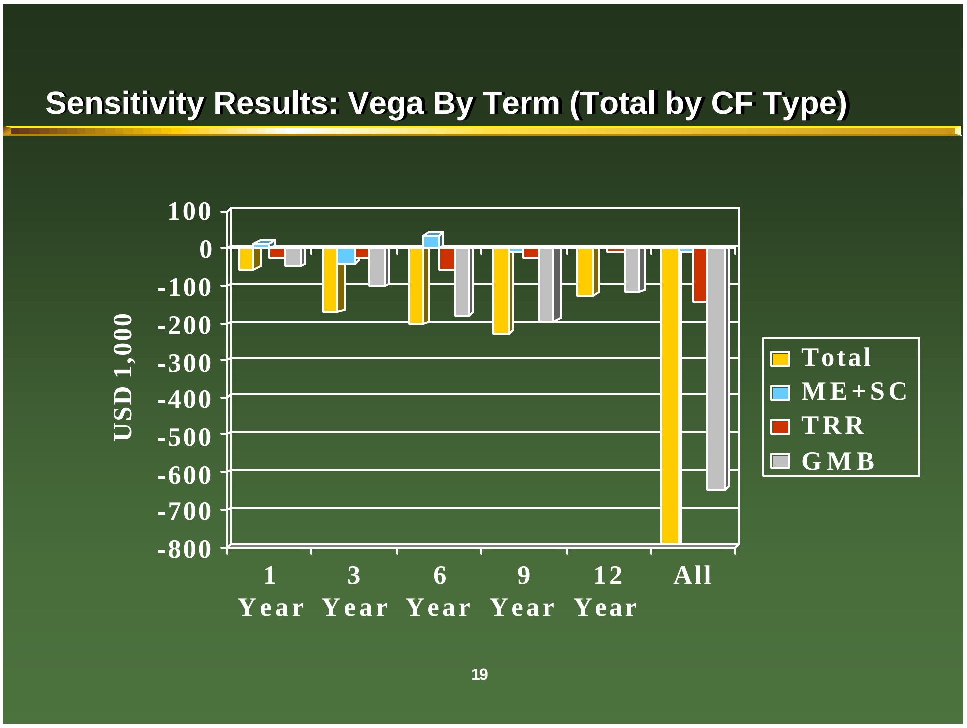

Chart 10

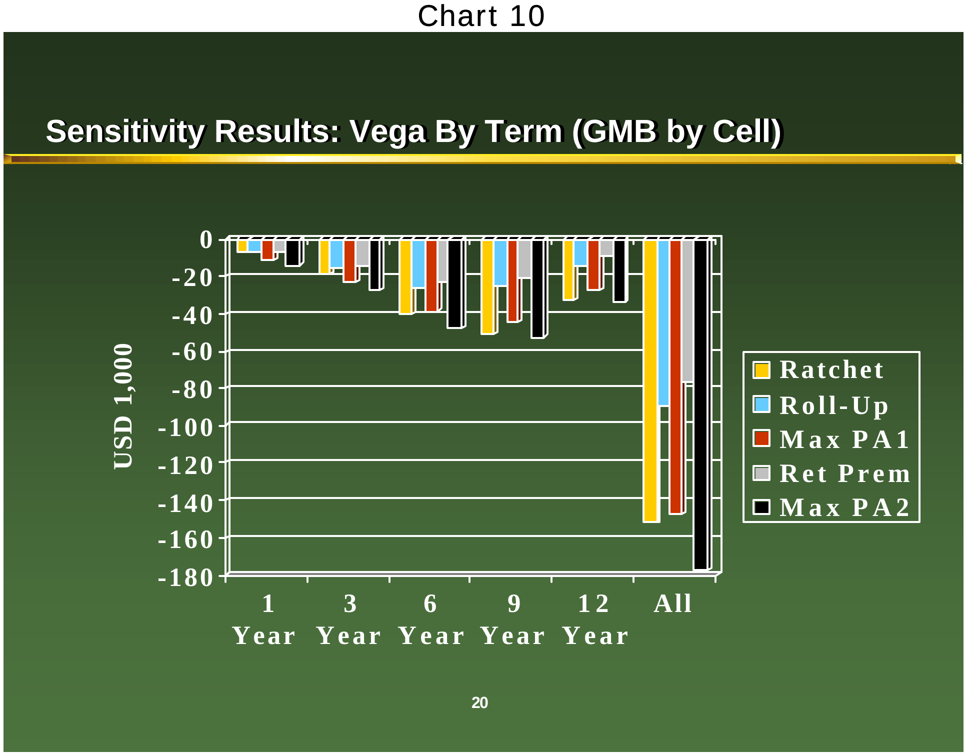

Chart 11

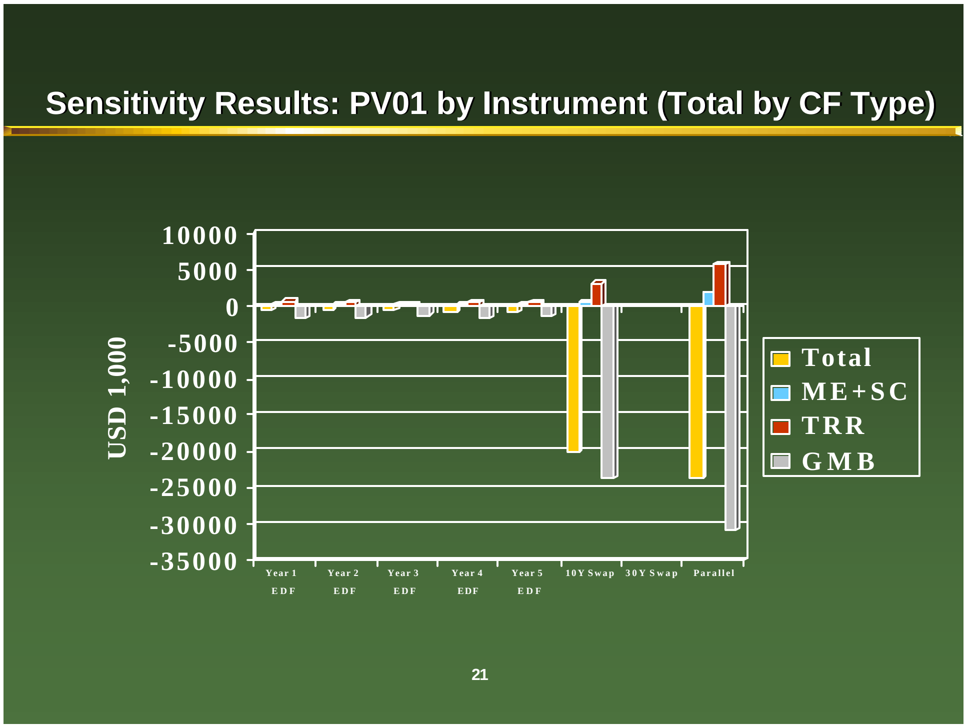

Chart 12

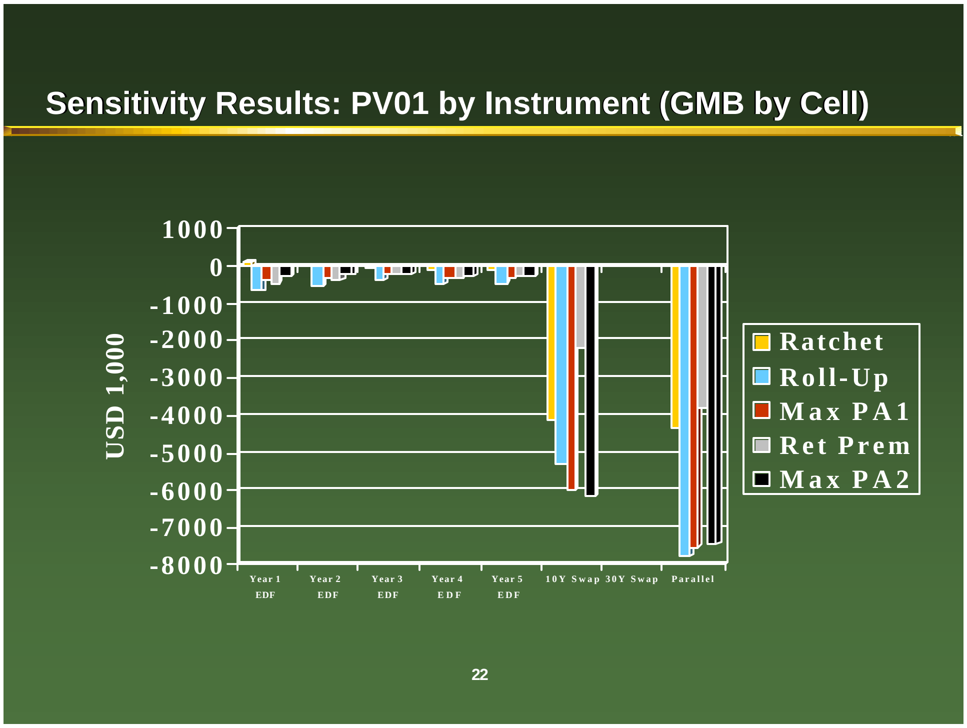

Chart 13

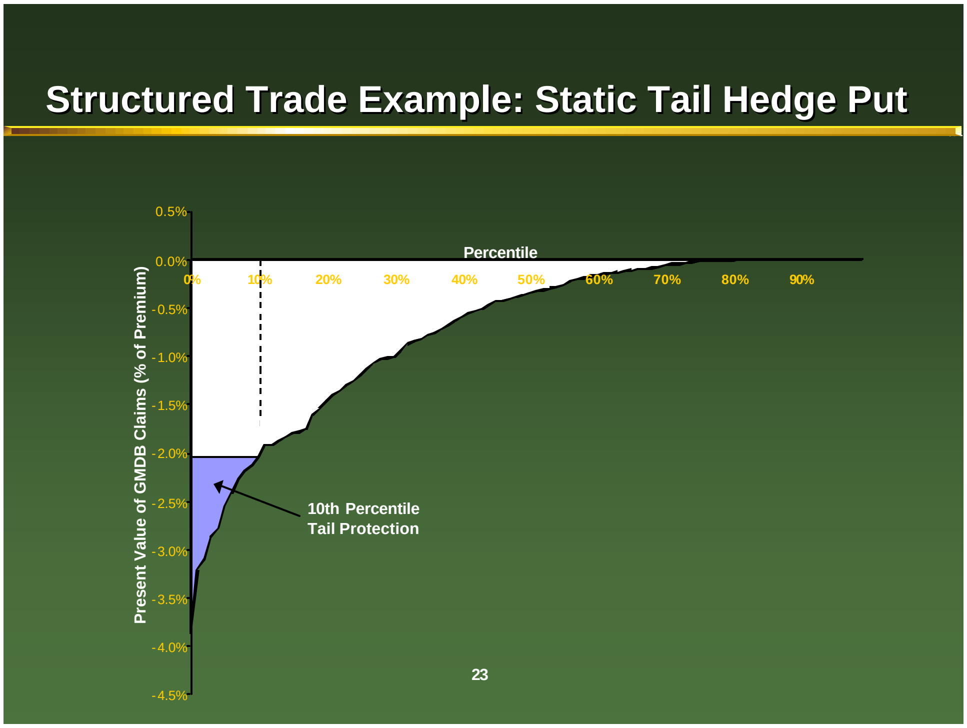

Chart 14

Scenario Optimization-continued

$$e_s = \sum_{i=1}^{n} [w_i \cdot P_H^i(s)] - P_T(s)$$

$$R_N = \sum_{s=1}^{S} p_s \, |\min(e_s, 0)|$$

Chart 15

Scenario Optimization-continued

$$e_N = \min(e_1, e_2, \cdots, e_S)$$

$$\text{Minimize}(\ R_N\ )$$

$$e_N > a \quad \text{and} \quad \sum_{i=1}^{n} w_i V_H^i < b$$

Chart 16

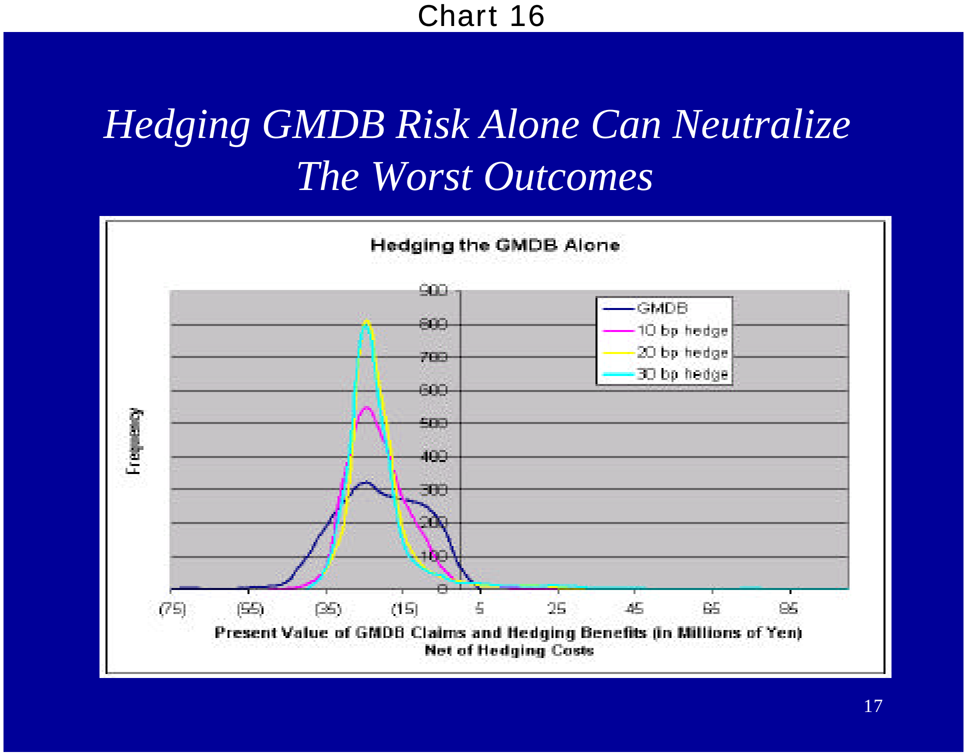

Chart 17

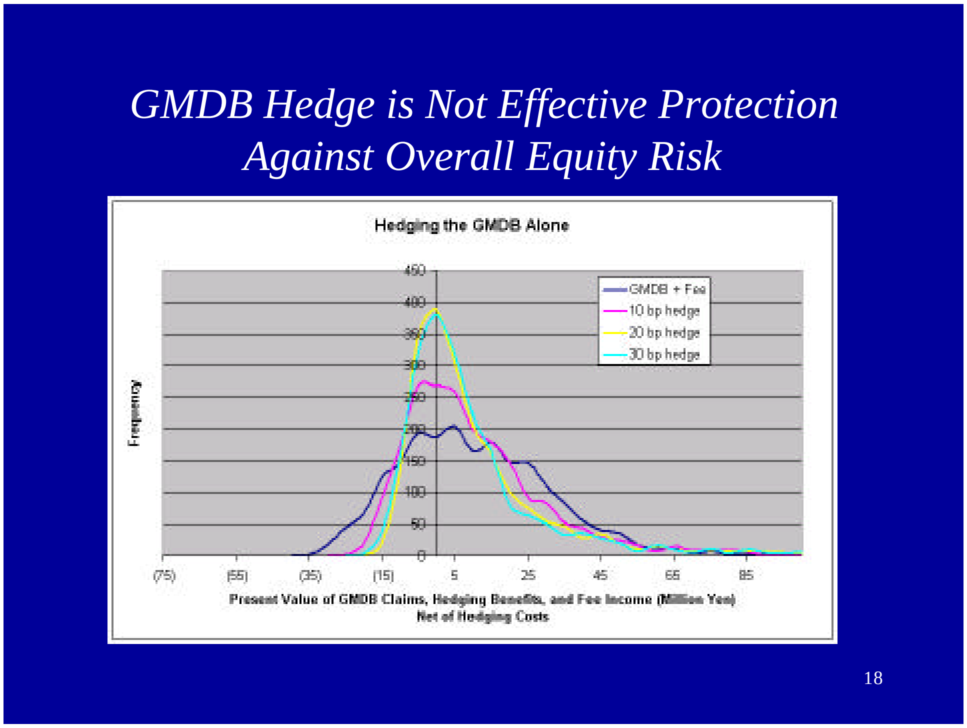

Chart 18

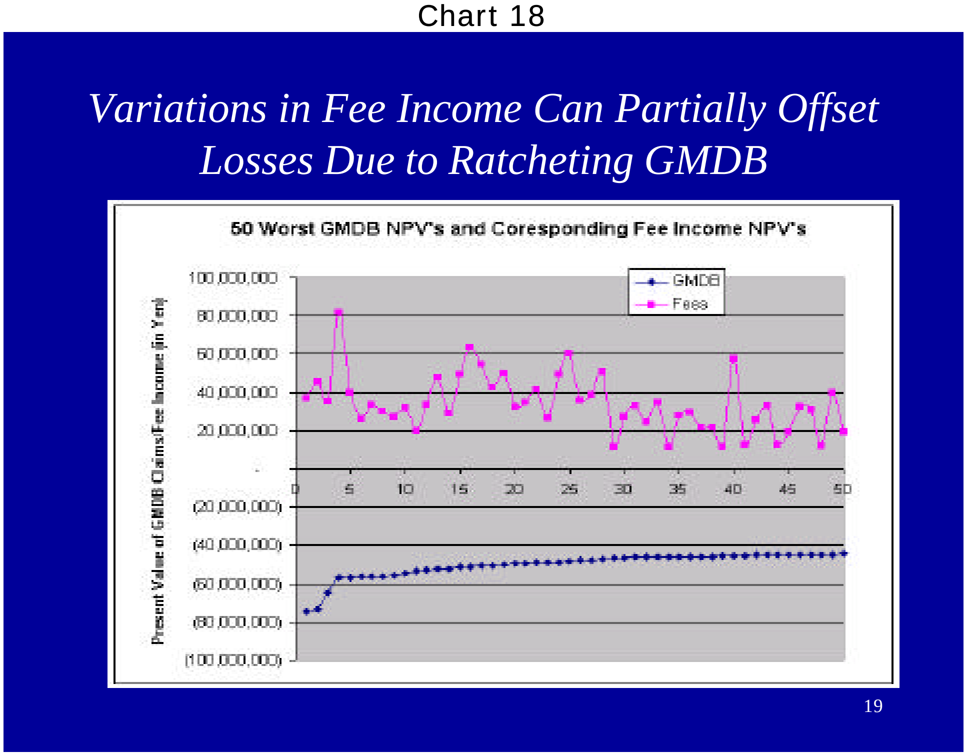

Chart 19

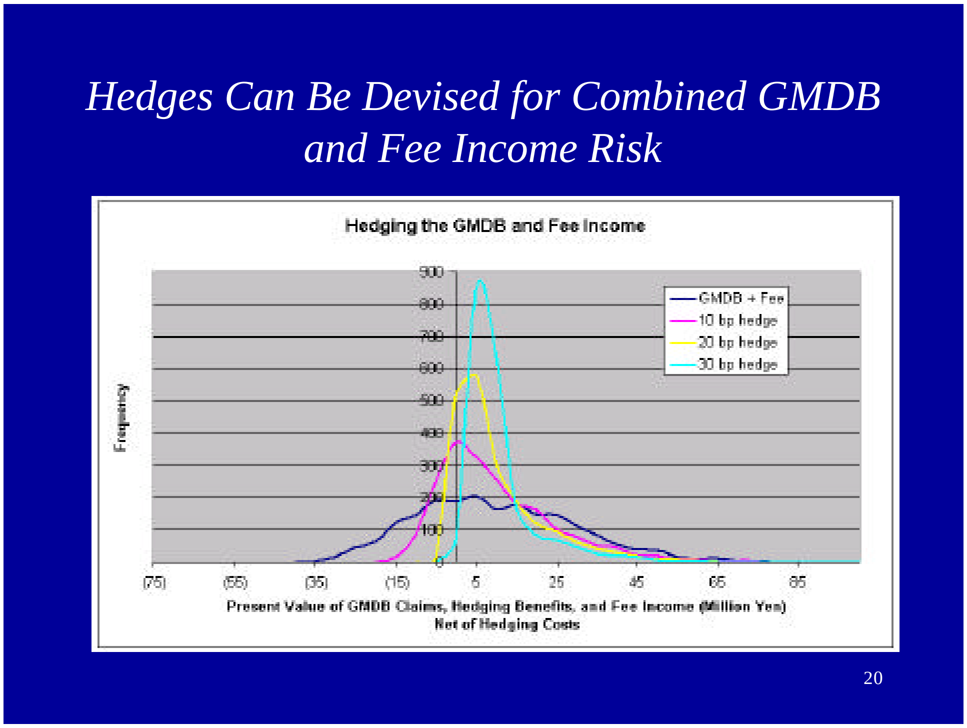

Chart 20

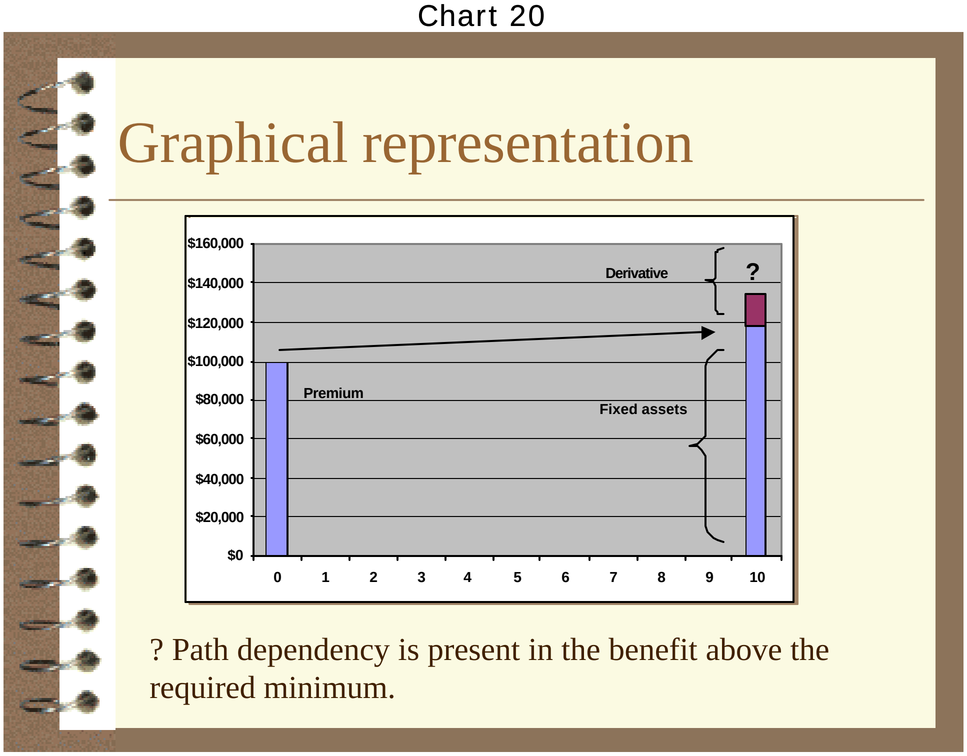

Chart 21

Hedge performance - Naked

- P&L at year 10 = Premium * 1.05 ^ 10 - Premium * 0.875 * 1.03 ^ 10 + ????, where ???? is positive.

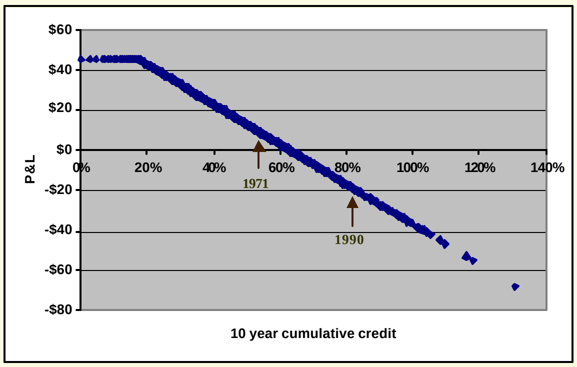

Historically a poor strategy, failing in all but two of the last 30 years

Standard deviation of results > 120% of the average gain

Chart 22

Hedge performance - hedge above minimum required growth

- Strategy
  - Purchase call based on implied guaranteed growth rate to end of the term
  - At issue the guarantee of 118% of premium over 10 years, implies 1.67% strike above current level.
  - Credit of 5% in year 1, implies 1.18/1.05 ^ (1/9) - 1 = 1.31% strike in second year, etc...

Chart 23

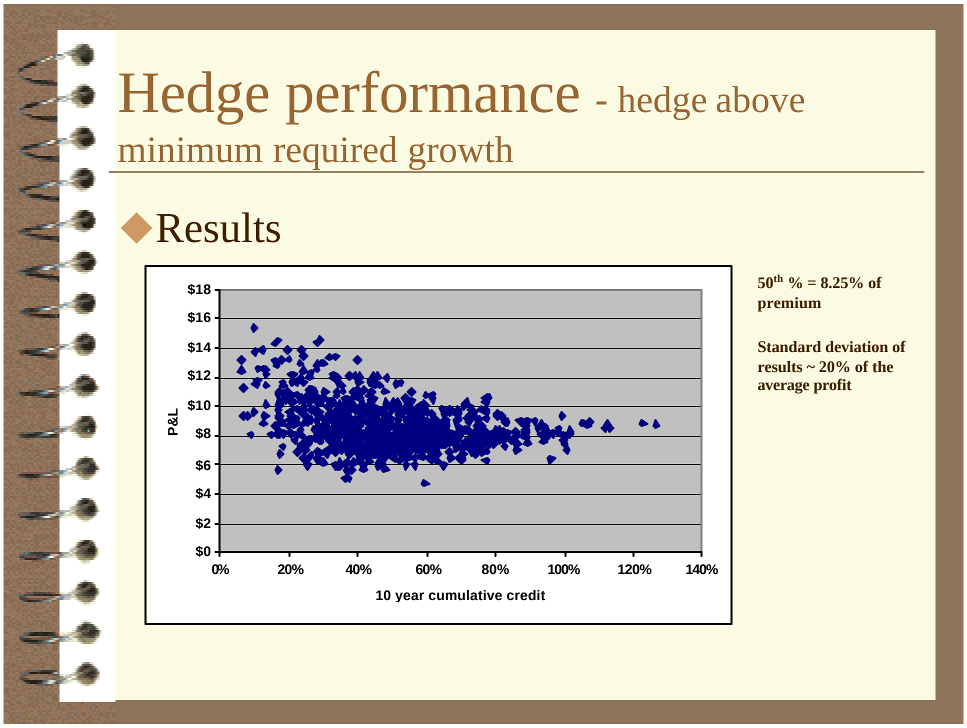

Chart 24

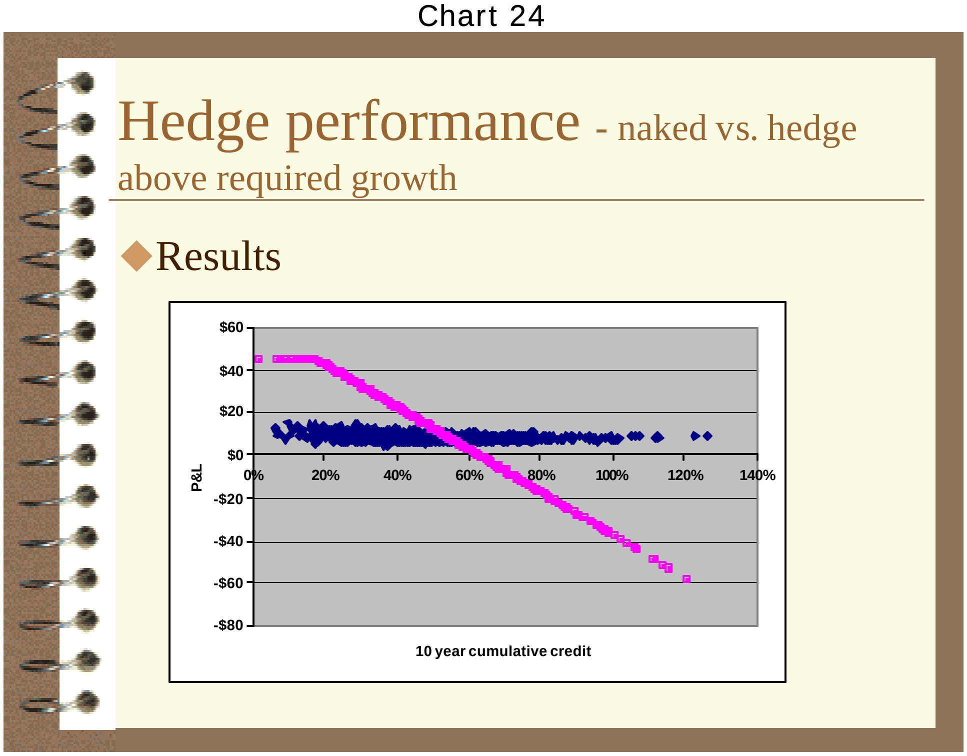

Chart 25

Hedge performance - hedge your option budget

- Results

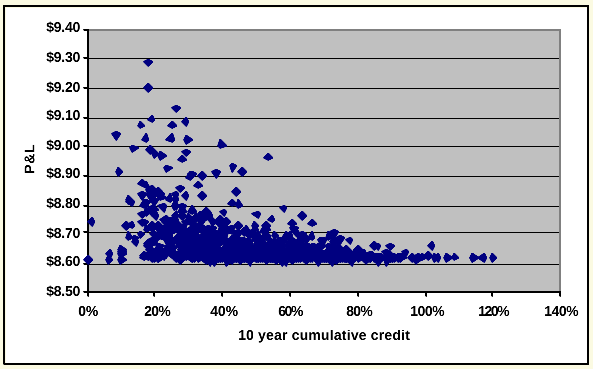

50th % = 8.64% of premium

Standard deviation of results < 1% of the average profit

Chart 26

Hedge performance - hedge your option budget vs. hedge above guarantee minimum

- Results

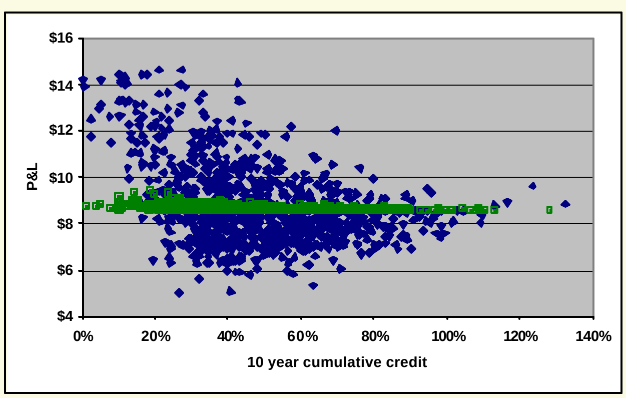

Chart 27

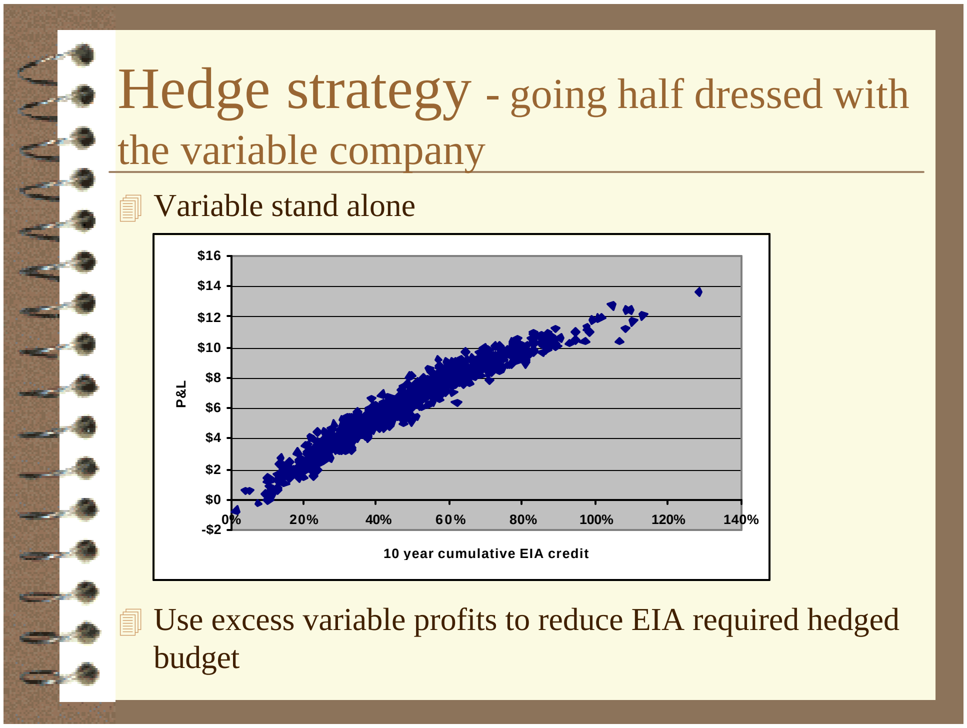

Chart 28