RECORD, Volume 27, No 1*

Colorado Springs Spring Meeting
May 30-31, 2002

Session 64PD
Dynamic Hedging

Track: Product Development/Investment

Moderator: MARSHALL C. GREENBAUM
Panelist: K. RAVINDRAN†

Summary: Capital market instruments are being used to hedge the risks assumed by writers of variable life insurance, variable annuities and equity-indexed products. Instructors use case studies of benefits and guarantees similar to those available in the marketplace to illustrate the effectiveness of the various hedges. Attendees learn the nature of the risks associated with equity-based and equity-linked products, and how those products can be managed with the use of capital market instruments.

MR. MARSHALL C. GREENBAUM: I'm with Constellation Financial Management Company. For those of you not familiar with my company, the easiest way to describe it is we purchase deferred acquisition cost (DAC) and in a manner that eliminates it from a company's balance sheet. All of our financing transactions are completely non-recourse. A typical transaction involves advancing an amount equivalent to a broker's commission and then collecting variable (market-driven) fees (such as a 12b-1, a distribution fee charges as a percentage of a fund's net asset value) for a defined period of time. We accept all of the risk associated with the variable fees.

We are fortunate to be joined by Mr. Ravindran, who has experience helping insurers with dynamic hedging issues. He is the founder of Annuity Systems, Inc., a software and consulting company dedicated to the risk management of annuity products. I'm going to focus on dynamic hedging as it relates to variable annuities. Mr. Ravindran is going to gear his presentation to equity index annuities.

†Mr. K. Ravindran, not a member of the sponsoring organizations, is president at Annuity Systems Inc. in Toronto, Ontario.

There are a couple of main points that I'm going to make throughout my presentation that I want to state up-front. First, the variable annuity (VA) marketplace and the issuance of these products are going to be drastically changing in the near future. Risk-based capital (RBC) requirements for variable products with guarantees are going to be drastically increased. I'm going to give a brief overview of what is pending and the overall impact it will have on the variable product industry.

Another main point I'm going to make throughout my presentation is there has been much focus on hedging as it relates to guaranteed minimum death and income benefits (GMDB and GMIB). For those of you that were able to attend Session 6PD, "Managing Equity Guarantees," Thomas Ho also stressed the point that there is an additional "animal in the room," the market risk associated with the revenue fees that are attached to these products.

Typically, mortality and expense (M&E) fees are charged against account value; the market goes up and more revenue is generated in dollar terms. When the market goes down, less fee revenue is collected. Another related point is that the actual sale of the variable fees is a risk management option. If you have the ability to fully collateralize a known amount at the sale of the product for future unknown fees to be collected, you can win or lose with the benefit of hindsight. The actual fees on a present value basis could have turned out to be more or less than the lump sum received upfront depending on how the market performed.

The third point is that dynamic hedging works. However, there are a lot of practical issues involved with dynamic hedging. In real practice, a dynamic hedger will assume "basis risk" and "market gap risk" just to name a couple. The market gap risk issue is that when you're trying to replicate a guaranteed minimum benefit with other financial instruments, you need to rely upon buying and selling those instruments. If there's a time when you can't buy or sell, you are going to have a gain/(loss) impact from your hedging program.

Because of the changing RBC requirement, hedging of variable products will be heavily utilized.

My presentation focuses on capital markets risk transfer. If you're going to use a dynamic hedging program, that's what you're looking to do; you're looking to take out the capital markets risk. It is difficult, although not impossible, to transfer other material risks, such as lapse and mortality risk through securitization.

Embedded in a variable annuity contract are two essential components of market risk. There's the guaranteed benefit option, but there are also the M&E fees collected. Additionally, sometimes there are surrender charges that are calculated against account value that fluctuates with the market. It is exactly like M&E fees in that if the account value goes down, you're going to collect a lower surrender charge in dollars. If the account value goes up, you're going to collect more in surrender charges. In a declining equity market, both pieces of the equation will be negatively impacted. The M&E fees (the assets) are lower and the guaranteed minimum benefits (GMBs), the liabilities, are increasing. A hedging metric should not only include the guarantees but be more holistic and include the M&E fees. Something like the M&E fee less the guaranteed benefits. Other impacts on earnings driven by the market, such as surrender charges or market-related expenses would be worthy of inclusion as well.

A select few writers are hedging with respect to the guaranteed options only. One of the main reasons that insurers have only hedged the options is that under current regulations there is not incentive to hedge the revenue risk. There is no capital or reserve relief for doing so. This is changing. If proposed regulations get enacted, insurers will get a capital benefit from hedging the M&E fee revenue as well as the guarantees.

Another reason hedging the M&E fee risk has been nonexistent is that the future M&E fee revenue expected to be collected over time is not collateralized on an insurer's balance sheet as an asset on a market value basis. They get played out over time as a cash flow event. After applying a hedge to the fees, under Financial Accounting Standard (FAS) 133, income statement volatility would be introduced. The hedge would be marked to market on the balance statement, but your M&E fees would not be, thus creating income statement volatility.

What are these pending capital regulations? Targeting year-end 2003, the American Academy of Actuaries has a sub committee on the NAIC life risk based capital working group working on the proposed regulations. They are indicating something similar to what is in place under Canadian regulations and in particular conditional tail expectation (CTE) calculations.

The CTE calculation works from a Monte Carlo simulation of cash flows. For each scenario, calculate the present value of your fees less guaranteed claims and average the worst 5th percentile of the present value figures (CTE95). This calculation would theoretically have the insurer hold claims to cover somewhere between the 95th and 100th percentile of all scenarios.

The U.S. committee is suggesting a CTE90 measure, differing from the Canadian calculation in a couple of key respects. The Canadian regulation is all cash flow based. In the United States, they are suggesting a statutory loss basis inclusive of expenses and taxes. The Monte Carlo scenarios would be generated in a manner similar to the Canadian calculation, a "real world" simulation process or projection of returns with assumed risk premiums calibrated to historical data. The scenarios need to satisfy calibration requirements insuring "fat tails" to the distribution of returns.

Another difference in the calculation is not having the ability to include a negative loss offsetting the calculation in the worst 5^{th} percentile. If one of those scenarios in the worst 5^{th} percentile is positive, you have to include that in the average as zero. My understanding in the Canadian calculation is that a positive result can offset the negative scenarios.

The other potential divergence from the Canadian calculation is the concept of holding the maximum present value loss in any scenario over all years in the projection. You can't allow some sort of recovery in the market in the end years of the projection to lower the loss in a particular scenario. This is dramatically increasing the calculation and is being scrutinized for appropriateness.

The proposed regulations will substantially increase RBC requirements for variable products. Today, insurers are holding somewhere between 50-100 basis point of account value, and according to my calculations, may need to be increased 5-10 times that amount depending on the richness of the guarantee benefits offered.

The committee is recognizing the ability to offset the calculation with a hedge credit, although they have not given guidance to the extent of the credit offset. Monte Carlo results will be evaluated with and without the hedge program in place to drive the maximum amount of potential credit. A percentage of that calculated amount may be applied. Insurers will under the proposed regulation have a significant incentive to hedge their risk. Hopefully the hedge offset credit will be substantial if full credit is not giving further encouragement to hedge.

What are the different kinds of hedging alternatives currently available? Certainly, going naked is not a hedge alternative, but it's a baseline for comparison to hedging alternatives. If you're not hedging the market risks, then it would be prudent to hold capital against potential claims.

In the past, an alternative to doing nothing was GMDB reinsurance. They are usually structured to protect against the guaranteed benefits embedded in the policy only. The status of this market is always changing from one to two writers; it’s an extremely thin if not nonexistent market.

Another alternative is a total return swap. This would be structured against the M&E revenue. One of the nice aspects about a total return swap is that it is FAS 133 friendly. A long-term contract, say 8-10 years of fixed cash flows for the actual M&E cash flows received is a perfect hedge. It's a perfect offset. There's no "hedging ineffectiveness." You avoid undesirable income statement volatility.

There are also static option alternatives, the purchase of long dated put options and/or the sale of call options. The concept being that after purchase they would be hold-to-maturity instruments for capital calculation and transactional cost purposes. This will be the viable market in the future for hedging the VA risk. These long-dated exotic contracts will match very closely the embedded risks in the variable annuity products. They will also serve to maximize the capital relief in the CTE calculation methodology. Having the contract payoffs based on the subaccount returns embedded in the policy and not index returns will allow for optimal accounting presentation. They will include full risk transfer of the market risk inclusive of the basis risk.

Another hedge alternative is dynamic hedging. This involves the use of short-dated derivatives instruments held on the balance sheet to hedge against undesirable losses.

Yet another alternative is the "sale of revenue fees." Basically, take your future revenue fees and simply sell it. I'm always reminded of a quote from the current CEO of Merrill Lynch who said a few years back after their bond desk experienced some hedge losses, "The only perfect hedge is when somebody else owns it." He was referring to something similar to dynamic hedging with index trades, where the hedge portfolio did not exactly match the underlying.

What is dynamic hedging? You have some sort of put option that you want to hedge against. You calculate your hedge ratio and you enter into a futures contract intended to replicate the put option sensitivity. Then you step forward in time and recalculate that put option, if the exposure on your asset side (the futures contract) no longer matches the put option, you need to rebalance the futures contract. Essentially, you have some liability that you want to hedge against. You can't actually (or do not want to) sell the exposure exactly, so you use something very closely related to give you a very highly correlated offsetting exposure.

After calculating the sensitivity to the exposure that you're trying to protect against, create the return profile that you need. For example, compare short index features against GMBD exposure. You will need to constantly readjust the short position depending on market movements. Eventually, if and when a GMDB cash flow occurs, the net results of the futures trade will pay the claim.

Dynamic hedging relies on liquid and reasonably continuous markets with low or moderate transaction costs. If there's a period of time when you can't transact, there is going to be a P&L event from the hedge program. Very large transaction costs will also drive up the costs of the program and can in many instances make the program cost prohibitive. It will eat away at the product's economics. The impact will be played out over time. It's always prudent when you're thinking about a dynamic hedging program, to actually model in the transaction costs and all of the practical implications of dynamic hedging to derive at the optimal strategy.

Again, there are two main pieces to the hedging metric for a variable product: the M&E piece and the GMBD piece. The M&E piece is a linear risk exposure. It's symmetrical—when the market goes up you win; when the market goes down you lose. It's the exact profile of a stock or fund. An insurer with an unhedged VA exposure has essentially provided its shareholders with an investment in the actual subaccounts underlying the variable annuity product. Is the VA product line the optimal way for shareholders to get equity market exposure? The M&E fees are the easier piece of the hedge metric equation to hedge. Because it's linear, it only contains a *delta* exposure, with no higher order *Greek* letters to worry about. Delta is the sensitivity of the option to the underlying itself. As a result, you can delta hedge the M&E exposure with liquidly traded futures contracts that do not have to be monitored as frequently as the GMBs. That's a really nice aspect of hedging M&E fees.

The guaranteed benefits piece of the hedging metric is much more complicated, because you've got to constantly monitor it. The exposure is not linear—when the market goes up you don't have a claim; when the market goes down you have a claim. It's your typical option payoff profile. Options have gamma exposure, or sensitivity with respect to delta exposure. Gamma is what's causing you to have to frequently monitor your hedge positions. When the market goes up, your delta's going to go to zero on your guaranteed benefits. You can sell out of the hedges and be done with hedges eventually. When the market goes up, your option value approaches zero, which is referred to as out-of-the-money. You're not going to need any hedge instrument anymore. As the market starts to go down, you need to eventually be one-for-one hedged against the GMDB. When the GMBD is sufficiently in-the-money, it's going to have an exact movement with the underlying and you need to be fully hedged against it. Thus, the management of the guaranteed benefit benefits requires much more frequent rebalancing, monitoring along with the sophisticated infrastructure to accommodate a dynamic hedge program.

The essence of dynamic hedging is that you need to do a bunch of fair market value calculations. Once you've decided on the hedging metric that you're going to use, whether it's just the M&E or just the GMBD or it's a combination of the two, you need to basically do a bunch of fair market value calculations to get your hands around the asset/liability. I'm not going to get into the details of how that fair value calculation is done. There's plenty of good literature detailing the process of option pricing. In a nutshell, you need to utilize a risk-neutral not "real world" framework.

Because the GMBs are path-dependent, unfortunately we can't rely upon closed-form solutions for the fair market value calculations. The ultimate payoffs are dependent on not only the final values but how the market got to those values. For instance, the behavior of the policyholders is market dependent. Thus we don't know how many policies will persist by just knowing where the stock market is at a point in time; we need to also know how it got there. In a bad market, policyholders may have lapsed along the way only to see a market rebound. Thus the fair market calculations rely upon Monte Carlo simulation. Once you have a defined hedge metric, you need to calculate your sensitivities to capital market parameters, the sensitivity to the equity markets (delta) at a minimum. If you want hedge against other higher order Greeks, you need to calculate them as well.

When calculating the fair market values, you need to adjust not only the current market parameters but the policyholder information as well. You need to know who's lapsing, who's still there, how their asset allocation has changed, etc. It's a constant monitoring process, not only of the capital markets but also of the actual policyholders remaining on the books.

Once you have your liability calculations done along with their sensitivities to market parameters, you need to take a look at your hedge portfolio of derivative instruments. If you've just started the process, you need to enter into contracts that exactly offset your delta exposure. Calculating the Greek exposure of the derivative portfolio is a bit easier than the liabilities because closed-form solutions can typically be relied upon without Monte Carlo simulations. They may not be such simple formulas, but at least they're straightforward to program and process. There are numerous software packages that can be relied upon for this process.

Once you have the Greek exposures on the liability side and you have the Greek exposures on the hedge side, then you simply make sure that they are aligned. There are probably some real subtle things that you could do on your liability to accomplish this, but realigning them will mainly be accomplished by rebalancing the derivative portfolio. Once it's rebalanced, after any period of time has elapsed, you need to look at it again to make sure they are still aligned within an acceptable tolerance. It's a constant monitoring process.

There are several limitations to doing dynamic hedging. I alluded to them earlier. One of the key things is that you really need to take a careful look at when you hedge the higher order Greeks like gamma and vega is that it's very easy to let transaction costs eat away at your profits. The bid-ask spreads in the option marketplace are substantial compared to the futures market. If the options traded to offset the gamma/vegas exposure need to be traded often, the bid-ask spread will on option hedges will create substantial costs to the program. Another major item of concern is "basis risk." Basis risk occurs when you realistically look to hedge your sub accounts, exposure to actively managed funds, that don't really match an index exactly. Because of the liquid market trading condition of dynamic hedging, you become limited to using a limited basket of indices to hedge with.
Even after using, say, eight or nine indices, you're left with substantial basis risk. It doesn't matter how much volume you have, it's how many funds and are they actively managed? Even 30 funds will not match a basket of indices very well. Most variable annuity writers will still retain basis risk volatility of 4-10 percent after implementing a dynamic hedge program.

Let's say you have a $10,000 investment in a variable annuity and that the guaranteed benefit option cost (fair market value) is two percent of the deposit, or $200. Let's assume five percent basis risk volatility at two standard deviations, or 10 percent of $200 is $20. Thus, you can have a P&L event of plus or minus at least $20 on an annualized basis with a five percent probability. To put that in relation to the income of the VA policy, if you're earning 100 basis points M&E, that's $100 on your $10,000 investment. So you're looking at a $100 M&E investment with a plus or minus $20 impact from your hedging program. That's not insignificant. This can result even after implementing a dynamic hedging program with a substantial investment in infrastructure including the proper controls to effectively execute.

Let's pretend we're in the year-end 2003 and the proposed regulations for VAs are in effect. I've implemented a model for statutory earnings based on the Monte Carlo simulations approach. On each one of the paths, the CTE90 calculation will prevail for determining statutory capital requirements dictating rating agency requirements via a multiple of the statutory requirement. At each projection point in time on a given scenario, we calculate the CTE90 result, a Monte Carlo valuation itself. Thus we have Monte Carlo valuations going on within a Monte Carlo valuation, a computationally intensive process. Once we have the simulation run, we can look at after-tax statutory distributable earnings as a proxy for shareholder value.

If we hold capital at CTE90, that is akin to an adjusted return-on-capital pricing approach. At each point in time we're holding a sufficient amount of capital to cover claims. Thus, after-tax distributable earnings holding CTE90 capital is a decent risk-adjusted return measure. After we run a number of experiments, each containing a different hedge alternative, we can compare via this return measure. It may not be the optimal risk-adjusted measure, maybe a CTE98 calculation is what your company would like to see but nonetheless it will provide guidance for the different hedge experiments.

The case study projects the earnings of a generic VA product. The average M&E fee is 140 basis points. The investment management fees are 95 basis points. Both fees are assessed against account value. A GMDB of premiums accumulated at five percent interest is included for an additional 25 basis of account value charge. For the capital market assumption, the individual policyholder accounts are assumed to return 11 percent annually with a volatility of 18 percent. Let's assume that the company holds 150 percent of that total balance sheet requirement that's calculated via the CTE 90 approach and that the company has a hurdle rate of 12 percent.

I limited the experiments to four. Obviously, there are an infinite number of experiments that can be run. The experiments performed were 1) just don't hedge, 2) reinsure your GMDB for 100 percent quota share for a charge of 40 basis points (In today's market it's not plausible, but was performed for comparison purposes. The 40 basis point charge is intended to be representative. Most reinsurance treaties executed in the past have put some sort of tail risk back on the direct writer. The 40 basis points is intended to be indicative of full GMBD risk transfer.) 3) a simplified dynamic hedging experiment of delta hedging the GMBD exposure only with futures contracts only. (Insurers initially would be inclined to focus on the GMDB as a first step even though the macro equation is what I am preaching. I have assumed that the insurer would retain a five percent residual volatility for this hedge program.) 4) a straight sale of fees of 75 basis point of account value for seven years, and a portion of the surrender charge schedule, 5.1 percent grading down to one percent. The insurer would receive up front an advance of 4.1 percent of deposits that would be immediately income recognizable. There would be no debt consequences to it. That is the key to getting the optimal economics out of this proposition.

Note, one of the major lynchpins to why we haven't seen many securitizations of fees has been the accounting. The sale of fees is very prominent in the mutual fund industry, where mutual fund companies finance their back-end loaded B share class of their funds. There hasn't been any, what I call "true sale" transactions, done with insurance products. At Constellation, we have recently put together a structure that gets the same accounting treatment as the mutual fund industry achieves, "true sale non-recourse financing" (i.e. no DAC on the books).

Let's look at the results and the risk-adjusted returns. Because they are all risk-adjusted returns, we can directly compare them to each other to see which experiment is optimal. For each scenario, I've calculated the present value of earnings at 12 percent, the company's hurdle rate, and then taken the average over all scenarios. If you do nothing, this product is priced to earn over 12 percent, even given the new capital proposed regulations. The internal rate of return (IRR) is around 16-17 percent resulting in a positive PVDE discounting at 12 percent. These IRRs would be more consistent with ones on fixed annuities. When holding minimal capital, VAs price with extremely high IRRs, which is somewhat insignificant, because the high IRRs do not necessarily turn into income in terms of dollars. You could think of the PVDE measure as shareholder value added over your 12 percent hurdle rate. The naked result is showing a positive result, but there's another key consideration here. The initial distributable earnings strain is 5.9 percent, in excess of fixed annuities. Capital consideration starts at 5.9 percent. The market goes down, and you do another CTE calculation, the capital requirement can go up to 8-10 percent, or even higher. If you look at that from a realistic standpoint, you have to be able to manage a very volatile capital situation and potentially raise capital at perhaps the absolute worst times (after markets have declined) when capital is a scarcity.

In the reinsurance experiment, recall that we have completely reinsured the GMBD. With respect to the initial distributable earning strain, since the guarantee risk is transferred, and we get the hedge offset in the CTE calculation, there only is an initial strain of 2.8 percent of initial premium. On a risk-adjusted basis, this result is worse than the "no hedge" experiment. This is a result of the reinsurance cost. With no reinsurance cost, we would be much better off. The economics of this experiment becomes solely a function of reinsurance cost. At a 40 basis point annualized charge, it destroys value. The reinsurance cost outweighed the capital relief benefit. However, we've solved the capital volatility problem to a large extent. Given a sole choice of reinsure or do nothing, I think many writers would opt for the reinsurance alternative because of the capital volatility issue.

The dynamic hedging result is a "great" result and may seem counter-intuitive at first glance. The $8,000 result goes up to almost $15,000 based on $1 million in premium. In this experiment, we dynamically hedge the market risk and accept the basis risk of our funds not matching the index. We reflect the program in the CTE capital calculations and have assumed 100 percent credit for the hedge offsets in the Monte Carlo results. If 100 percent credit is not achievable, the results may be somewhat overstated. In this situation, we have transferred all the market risk, we have considerably alleviated our capital problem and we earn a higher return. It's a "win, win, win" situation. Without CTE capital requirements, you can have a situation where with hedging, it appears that you give up much upside potential, particularly if you assume the equity markets perform well on average. With CTE capital calculations, it's exactly the opposite. You are better off economically with hedging in most circumstances. The reason is the benefit from not holding as much capital with the high cost associated with it (a 12 percent hurdle rate) which outweighs the upside give-up from hedging the cash flows. Even with the market returning 10 percent per annum, the result with hedging is eye opening. The insurer is much better off than without hedging even when no actual GMDB claims come to fruition.

In the "sale of fees" experiment, instead of setting up the daunting dynamic hedging infrastructure that is necessary, we just enter into some legal "true sale" contracts. You get the same economics out of it as the hedge result with a lot less hassle. There is a difference with respect to the capital exposure. We reduced it initially to four percent, but since we have not hedged the GMBD, we still have some capital volatility issues to contend with. From a pure economic standpoint, you are on equal terms with the hedged result. In a situation where future dynamic hedge insurers may be able to out price the others, this may be a way for some writers to remain competitive.

There are some accounting considerations with respect to dynamic hedge programs. For the experiment, on the statutory side, since we relied upon short-dated futures, the contracts turn into cash. For example, if you mark to market your GMBD option and say it's $20 million at time zero, all of a sudden, the market goes way up and the value of that option goes to zero. Since you've hedged and if everything works perfectly, you will have to make net payments of $20 million dollars offsetting the $20MM gain. The $20MM is a cash event and you have to pay that out now. You have hedged against something that's could have happened in the future. Thus, the program has created statutory income statement volatility. There is some offset from the AG 34 requirements, but the AG 34 isn't a mark to market calculation. It's not exactly symmetrical. You may get a little bit of an offset from it, but not a very good one.

On a GAAP basis, it's my understanding that one of the writers that has a program and is dynamically hedging is able to deem the hedges in the program a "fair value hedge," thus the marked to market changes of the hedges reflected in the income statement are very closely offset by the marked to market changes of the bifurcated liabilities (the GMDB option). Under FAS 133, any "hedging ineffectiveness" must be recognized immediately. The five percent residual volatility or basis risk will get reflected in the income statement as it occurs. That will create some degree of additional income statement volatility. However, it is mild in comparison to the overall capital market volatility that has been removed via the hedge program.

There are some hedge alternatives other than the ones discussed above. A total return swap is one. It is a symmetrical trade giving up upside for downside protection. The other way is to massage that swap a bit by adding a call option. The insurer will receive some upside participation by paying slightly more for the downside protection. For instance, in a straight swap, you can guarantee say five percent returns or you can guarantee four percent while retaining a percentage of returns in excess of five percent. Equity indexed annuities are typically structured in a similar manner.

Another twist to the straight swap is a swaption, the ability to enter into a swap if the market goes down. In this structure, you can enter into the swap at attractive turns if the markets go down. If the market goes up, you retain the upside and just give up the cost of the swaption. This will help protect the tail exposure to the market risk and aid in reducing CTE calculations.

Lastly, structured products will be developed to exactly attack the tail that drives the CTE calculations. Basically these structures will contain option baskets structured to match the exposure more exactly than plain vanilla option will do. Most guaranteed features contain averaging and lookbacks are exotic in nature to begin with even without the actuarial risks. These structured products will give the VA writer the ability to remove the capital markets risk from their products at various levels of protection desired.

In summary, I hope I've made my point by now that in the future we are going to see much more hedging of variable products than we see today. If not, insurers will be penalized by extreme capital volatility and DAC write-downs. Shareholders will eventually be scrutinizing insurance companies for risk management programs centered around variable products.

MR. K. RAVINDRAN: I'm going to switch gears slightly from what Marshall had talked about in that I will focus on the equity-indexed annuity (EIA) products. In the process, I will give you my personal thoughts about dynamic hedging and what it takes to make dynamic hedging work in practice.

As Marshall said, dynamic hedging has been around for a very long time. Despite its duration of existence, there is still a big gap between theory and practice. What complicates the matter more is that even among practitioners, there tends to be a difference in how dynamic hedging is being applied. Furthermore, reading books on dynamic hedging and naively applying to practice what one reads without any relevant hands-on experience is a sure recipe for blow-ups (as I have personally evidenced by being brought in as a "forensic auditor" in the aftermath).

Having said this, I'm going to start by talking about the types of risks embedded in an EIA product followed by a description of the risk management process and my thoughts on some of the pitfalls associated with implementing a risk-management program. Thereafter, I'll address the topic of risk transfers while describing both dynamic hedging and quasi-static replication technique (something that I have personally used quite effectively as a market maker to immunize effectively against market gaps and volatility) methods.

There are all sorts of risks that one takes in selling a variable annuity or equity based annuity product and examples of these include earnings volatility and capital volatility. Whatever the implications of taking on these risks, it is important to note that ALL of them are subject to some form of volatility in market movements. The big reason why reinsurance has always been an attractive proposition to direct writers is due to its ability to reduce the volatility in ALL the risks simultaneously. As such, a complete reinsurance program would allow you to transfer ALL of your risks without having to worry about the earnings or accounting/regulatory implications. However, since complete reinsurance is currently not available, direct writers are forced to look elsewhere for alternatives to their traditional approach to risk management. As a consequence, they are faced with the problem of managing the "actual" market related risks while trying to contain the volatility of both the accounting and regulatory implications—as these metrics do not move in tandem with actual market movements. Hence, in employing an alternative approach to reinsurance for managing its risks or even an incomplete reinsurance program, the direct writer needs to decide which of the objective functions (e.g. volatility in CTE 95, volatility in earnings, volatility in reserves, claims-related risks, volatility in fee income, volatility in GAAP, just to name a few) is the most important to focus on since having an optimal risk management solution for one objective function does not imply a global optimality for this solution across multiple objective functions.

So, what constitutes a sound risk-management process? This question can be better answered quite succinctly at a high level using the following building blocks:

a) Risk Quantification
   i. What base lapse assumptions should one use?
   ii. What policyholder behavior model should be used to capture dynamic lapses?
   iii. What should the growth rate and the volatility of funds be?
   iv. What should the correlation between the funds be?
   v. What mathematical/economic model is going to be used to obtain the distribution of future fund values?
   vi. How should inputs (e.g. growth rates, volatilities, correlations, lapse rates) be quantified and objectively updated?

b) Objective function Identification

   i. Based on the risk quantification in a), what are the risks that are tolerable and what needs to be hedged?
   ii. How much risk needs to be hedged?
   iii. How much premium is available to hedge the risks identified in i. and ii.?
   iv. What should the trading limits be so that once market/lapse/volatility changes breach these limits, hedging strategies are deployed to manage these risks?

c) Risk Monitoring
   i. How frequently are the risks going to be monitored?
   ii. How to update the inputs to the pricing-hedging model identified in the Risk Quantification stage while minimizing the operational risks.

d) Risk Managing
   i. Can the risks be monitored live?
   ii. What are the types of instruments (assets that will be purchased to manage the risks) that need to be purchased?
   iii. What controls (i.e. systems, legal, data, process) need to be in place?
   iv. How to time the transactions of the trades so as to reduce transactions costs?

e) Repeat the above steps again.

One other related question that still needs to be answered relates to types of risk-transfer methods that are currently available in the marketplace. There are basically three ways for transferring risks. These are:

a) Self-Insurance (not doing anything)
b) Reinsurance (on your identified objective functions)
c) Capital Markets Hedging (which includes everything from the most dynamic (i.e. futures only) to the most static (i.e. customized OTC long-dated options)) but executed using the OTC market or the exchange traded market

It is important to note that all these strategies are all intimately connected and implementing a dynamic hedging program doesn't mean that you cannot move to a reinsurance program. At the end of the day, as a risk manager, you want a risk management program that's quite complete and versatile in the sense that you want to be able to move from any one of strategy to another without incurring a lot of transactions costs, while ultimately getting you to a static hedge in the cheapest possible manner. This is akin to driving around a block in search of a parking spot that is closest to your destination without consuming too much gas going around the block in circles. In a similar vein, you would want to find a static hedge that is matches your risks exactly for the price that you want. Until you can get your static hedge, you want to be able to maintain a hedging program using instruments that can be easily liquidated or transformed into your ideal hedge.

Clearly, when self-insuring, there is no need for any rebalancing. As such, once this has been agreed upon (internally as a sound risk management strategy), you are only constantly monitoring your risks without doing any hedging. On the other extreme, if you have reinsured everything at a price that you like, then there's nothing to manage. Thus the more you get away from these two extremes, the more dynamic your hedging program becomes and the more frequently you have to rebalance, which translates directly to the amount of work and resources required to maintain a dynamic program which is analogous to some of the comments that Marshall had made about the amount of resources that are required to run a good dynamic hedging program.

So what is dynamic hedging? Unlike conventional wisdom, it has got nothing to do with using futures, forwards, swaps and/or options. It has got everything to do with how frequently you're running this whole program (i.e. daily, weekly, monthly or yearly). Succinctly put, the more frequently you do it, the more dynamic the program gets and the more you need to react. The less frequently you do it, the less dynamic (i.e. the more static) it gets. That's the whole paradigm underlying a dynamically hedged program.

With this in mind, I'm now going to go over a simple example of a dynamic hedging exercise. Since my talk is supposed to be focused on EIAs, I am going to talk about dynamically hedging a call option. Assume a 10-year stock call option, stock price of $100, strike price of $100, no dividends, four percent risk-free rate and an increasing volatility term structure from 16 percent to 25 percent (this is called an implied volatility term structure since it typically increases as you walk out in time).

Suppose that this call option has been sold for $45. How can one now manage the risks arising from the sale of this option? The first strategy for the writer of this call option is to self-insure. You collect the $45 premium and pray that 10 years later everything is all perfect.

The second alternative is to reinsure. In my example, I have used $40 as the reinsurance premium (although this doesn't have to be $40 and could be $48—in which case you're incurring a loss at the inception of the contract). By selling the option for $45, turning around and buying this reinsurance for $40, it would leave the writer with a tidy profit and no market exposure.

The third strategy revolves around hedging 100 percent of the risks upfront. To understand this better, first note that on the day of option expiry, should this option finish in the money, the buyer will exercise the option by buying the stock from writer at the strike price. Thus, in hedging 100 percent upfront, I was alluding to the fact that the writer would be buying the stock upfront for $100, warehouse it for 10 years and then transfer the stock if the option is exercised.

Unlike the third strategy, the fourth strategy involves only purchasing 50 percent of the stock upfront and the fifth strategy comprises one that is dynamically hedged. Now that I have quickly summarized these strategies, I would like to discuss them in further detail.

In strategy one, the writer is actually losing money as soon as the stock gets above $100 on maturity date since the buyer would exercise the option. The further the value of the stock increases, the more the writer is going to start losing. On the other hand, when the stock value is less than $100 on maturity date, the writer has collected the premium and has not incurred any loss. Having said this, I want to reiterate that taking on the risks naked and not doing anything can be a legitimate risk management strategy if you have a view that the stock value is not going to exceed $100 10 years later. Without such a view, taking on the risks naked is a different story altogether.

The second strategy, as you may recall, involves reinsurance. In this case, all your risks have been neutralized and you've just got a flat profit line regardless of the value of the stock at the end of 10 years (demonstrating a locked-in profit in excess of $5).

As per the third strategy, if you had bought the stock right up front, what would have happened? In the event the stock price exceeds $100 at the end of 10 years, the writer has to just deliver the stock and the maximum profit is capped out. On the other hand, if the converse happens, the writer has to liquidate the stock 10 years later. It should not come as any surprise that the fourth strategy (which involves a 50 percent upfront stock purchase) falls between the first and the third strategy.

The fifth strategy can be explained better by first assuming that hedging (or rebalancing) is only done once every two years during the 10-year life. Let me first start by purchasing 84 percent upfront (don't worry about how I got these numbers for the time being). Two years later the stock price goes down to $70.37 and I liquidate 27 percent of the stock so as to be holding 57 percent of stock. Following through with this, suppose that the stock prices every two years until the maturity of the options were $78.60 (Year 4), $114.10 (Year 6), $107.73 (Year 8) and $117.01 (Year 10), then the percentage of stock held at these times would be 58 percent, 85 percent, 78 percent and 100 percent respectively. So the whole idea here is that as the option goes into the money, the writer ends up holding 100 percent of the stock—in the process incurring $36.20 as the hedging cost. This cost is a function of the path (amongst other things, e.g. volatility, transaction cost, frequency of rebalancing etc). Thus, if you run 10,000 paths, you're going to have 10,000 associated costs—each corresponding to a path. Using this we can build a risk-profile curve associated with taking on these risks hedged—illustrating that by dynamically hedging the risks, the risk profile curve tends to flatten. This observation is true regardless of whether you are hedging the M&E fees or the 95 percent CTE values for the capital requirement and hence the benefits of hedging. So what exactly were we doing when we were dynamically hedging? To answer that question, first observe that a change in the liability value can be attributed to a

a) Change in the "underlying value" multiplied by "delta"
b) Change in the "underlying value squared" multiplied by "gamma" and a constant
c) Change in the "volatility" multiplied by "vega" etc.

where the "liability" mentioned above can be any objective function (e.g. claims, M&E revenue, 95 percent CTE). Furthermore, any changes in the "liability" can be immunized by neutralizing changes in its contributing drivers (e.g. delta, gamma, vega etc.). As a consequence, the greater (lesser) the number of contributing drivers neutralized, the lesser (greater) the changes in the "liability" and the more robust (dynamic) your hedging strategy. Having said this, the greater the number of contributing drivers neutralized, the more expensive the hedging program—although it is not necessarily true that to do this you would need a greater variety of instruments.

To hedge cheaply, one would need to immunize as little of the contributing drivers as possible, which as a consequence would imply frequently rebalancing the hedges. The catch, however, is that the greater the frequency of rebalancing, the greater the transaction costs, and hence the hedging costs associated with the program. On the other hand, although the transaction costs can be minimized by reducing the frequency of rebalancing, this reduction in frequency can lead to an increase in tracking error between the assets and the liabilities leading to an ineffective hedging program. The question then becomes one of balance. It is always the case that two companies offering exactly the same products and looking at risks similarly would have two different risk thresholds. Since the threshold point is highly dependent on how much risk a company can stomach on a day to day fluctuation or on a weekly fluctuation, it would be different across companies. When discussing our dynamic hedging example, one of the things that I had talked about was delta hedging. In practice, you have to set aside what I call a delta trading limit. If you hedge every minute of the day for any single market movement, your change in delta would become close to zero—in the process making the hedging program an expensive exercise. Thus, you would have to allow for some fluctuation in the delta (defined by the delta trading limit) and not hedge against every single market movement so that the hedging costs become more bearable. An extension to this idea is to also incorporate gamma limits, vega limits etc. Although the problem gets more complicated, all these things are crucial because the very moment something changes, you need to know what kind of trades you can execute in order to bring the risks back within the limits.

In practice, you monitor the net deltas, gammas, vegas, etc., arising from both the assets and liabilities. As long as you are still within the trading limits for each risk measure, you would be allowed to take a view on how to time the execution of the hedges. As soon as one of the trading limits is breached, you would not be allowed to take any more views and would be "forced" to execute the necessary trades to bring all the risks back in line within the prescribed limits. Some people also supplement these trading limits with value-at-risk (VAR) limits, total cumulative losses etc., so as to better contain their exposures.

Another important issue to consider when hedging is to be able to consistently hedge and report the results. More precisely, to calculate your delta, gamma, vega etc., for hedging purposes, one typically uses implied volatilities as inputs. However, when one is marking to market for reporting or accounting purposes, should implied or historical volatility be used? The inputs that are used for hedging have to be consistent with those used for reporting, otherwise you will end up hedging something different from what you are reporting (although the objective functions may be the same). In addition to these, one has to also analyze the impact of purchasing assets on the balance sheet (e.g. whether they can be treated as hedges, should they be treated as off-balance sheet items, what is the impact of capital after applying the hedging strategies) due to the fact that there could be a domino effect in implementing a nontraditional hedging program that can filter down to the balance sheet or earnings.

While hedging EIA products generally does not give rise to basis risks since the underlying supporting the products are indices, the same cannot be said for variable annuities. Since the underlying supporting any variable annuity tends to be actively managed funds and hedging can only be done using indices, this naturally gives rise to basis risks. So can basis risk be managed? Yes, it can be done, but with some work.

The ability to incorporate policyholder behavior for better risk management is also important. Working with a constant lapse assumption can only benefit you in a limited fashion. Obtaining data to incorporate surrender constraints, time lag from moment the notice is given to the moment the policy falls off the books, distribution channel operations, etc., allows you to fine tune your hedging program. As a consequence, you are able to better implement risk management strategies.

When hedging, it is very important for a risk manager to run an attribution analysis. Doing this would allow the risk manager to determine the drivers of the changes in M2M (e.g. changes in in force, lapses, market conditions, new businesses, etc). At the end of the day, to be able to perform a stellar risk management job, the risk manager should be able to attribute all the changes in M2M to any contributing factor that is subject to variation—so as to better focus on the culprit factors.

Another item to keep in mind when running a hedging program is to have the ability to report risk positions to senior management. Since people at the senior level don't have time to be able to sit down and filter through the mathematics that supports the analysis, it is important to illustrate this information via pictures and intuitive numbers. As a consequence, senior management can know exactly what risks they have on the books and whether they want to do anything about them. An example of this would be breaching the trading limits and still choosing not to execute any trades (since the risk-manager is confident about his/her view). By getting senior management involved and updated with what the company is exposed to, it is easier to empower the risk manager making crucial decisions related to rebalancing/taking a market view. Because of the condensing of information to pictures and intuitive numbers, it would be unreasonable for a risk manager to expect senior management to be as intimate as him or herself with the details or sources of these numbers.

When running a hedging program, you need to be able to quantify and manage your risks live. This is a very subjective statement since if the hedging program is a very static one (i.e. very robust and the assets tightly mirror the liabilities) then the need to quantify this live becomes unnecessary. On the other hand, if the risk management program is a more dynamic one, it is imperative for one to be able to quantify your risks live. In addition to being able to quantify and monitor your risks live, you also need the ability to time and execute trades effectively and efficiently. For example, it is sometimes the case that the yield curve or the volatility curve could be downward sloping, making it cheaper for you to lock in the "cheaper" part of the curve. Thus the ability to be flexible on the timing of the execution as opposed to executing trades on fixed dates of the month can make a lot of difference in terms of how much money you can save, sometimes even in the ballpark of $10-15 million, depending on the size of the trades.

Now that I have talked at length about dynamic hedging and some of practical issues associated with it, I would like now discuss quasi-static hedging. I personally have used this "Holy Grail" for many years in practice. The motivation behind this concept is a simple one in that you structure a robust (or "base") hedge that has very little fluctuations to market and assumption changes, although it would be robust 100 percent of the time. The power of this approach lies in the fact that you can customize your "base" hedge across the range of variables you want to take your view on (e.g. market not dropping more than 30 percent over the next month). Once the "base" position is put in place, you can then dynamically manage the residual risks (or noise) around the fluctuations on the "base" position.

Why is this approach better than dynamic hedging? One reason is due to the fact that it is less volatile and more predictable because the "base" component has been customized to be robust across scenarios you want to guard against. The second reason is that it can be potentially cheaper due to the fact that it would be a function of the execution time, market inconsistencies, and how the risks are sliced. The third reason for using a quasi-static approach is due to the static nature of the hedging strategy, in the process making this approach less resource intensive. The last reason for the pursuing a quasi-static approach lies in the fact that when you look at risk one way (i.e. the way the market looks at it), you'll find that it's very expensive to transfer the risks. However, if sliced and carved out differently to take a different form, it becomes cheaper to execute these slices—making the quasi-static approach a powerful one. Additionally, if it so happens that this "base" cannot be found at the inception of the liability, then one would first start off with a dynamic hedging program and then spot/lock opportunities walking through time. Thus, more you lock in, the more static the program gets.

What are the instruments are used for quasi-static hedging approach? Believe it or not it's a combination of self-insurance, reinsurance, capital markets—everything we have talked about. At the time of execution, one has to look at everything and quickly analyze where hedge prices are relative to one another. If a quasi-static hedging approach can be made to be cheap, static and effective, where is the catch? The first catch lies in the ability to cleverly slice and dice risk. In this manner, one can choose to transfer only components of risks that can be done at a favorable price. The second catch lies in the ability to know how to time the execution of the trades.

How does this "Holy Grail" apply to all the EIA products in the marketplace? As an example, consider a one-year monthly average option with a 100 percent participation rate subject to a maximum return of 10 percent. Purchasing the necessary averaging capped options would cost you around five percent. In employing the quasi-static hedging approach without taking on additional risks, the hedge cost can be brought down by about 35 basis points (as 35 basis points on a billion dollars can be quite a bit). This savings can be a lot bigger if you start introducing all the other finessing strategies that I have discussed thus far in my talk and can be as much as 1.5 percent.

The power of quasi-static hedging allows you to maintain the flexibility of the dynamic hedging program while providing you the robustness of the hedge. You don't want to be at the mercy of an investment bank and as such need to find niftier ways of trying to solve your hedging problems. The quasi-static approach helps you do exactly that. As mentioned earlier, if you are unable to execute a "base" quasi-static hedge, you would need to get into your dynamic hedging program on an interim basis until conditions become favorable—following which you would deploy the "base" hedge. Going to a full dynamic hedging program may not always be the optimal solution for everybody because of the related costs, infrastructure, etc. As such, quasi-static hedges provide a better alternative solution.

FROM THE FLOOR: You mentioned the CTE90. In your calculation did you actually use 250 percent of that? Are you suggesting that instead of RBC, we're using CTE90?

MR. GREENBAUM: I assumed that the insurer set up 150 percent of the CTE90 amount, actually the differential between CTE90 that total balance sheet requirement less reserves dictated by AG34. In addition, they also incorporated a five percent residual volatility estimate and incorporated that into the hedging program as modeled in the CTE90 calculation.

FROM THE FLOOR: In your calculation, when you added dynamic hedging, had you paid out 1,000 to 14,000, did you have any costs for transactions costs for your dynamic hedging?

MR. GREENBAUM: Yes, but because it strictly utilized futures contracts, the transaction costs involved are not material as opposed to programs that utilize option contracts. It didn't really have a large impact on the results.

FROM THE FLOOR: You're suggesting that you can slice and dice your risks in a manner to manage to beat the market potentially (or cut your costs). Some of it might be your assumptions or even your set of 1,000 or 10,000 expected future scenarios versus how the reinsurers view the world. If there are no differences, if their expectation of future is not different than ours and we're both using the same concepts, we should potentially come up with the same answers. What do we take for volatility in our stochastic analysis? Is the reason for these savings due to the fact one is using different inputs from the reinsurers?

To give an example, consider the mortality assumptions and your reserving calculations. It may be the case that you're assuming certain mortality and your reinsurers are assuming a lower mortality rate, making it possible for you to make money here. Isn't the saving of the money related to a variation in the stochasticity of the assumption sets?

MR. RAVINDRAN: No, not necessarily. Let's assume the worst case, that your assumptions are exactly the same. Let's also assume that you've got no view in terms of that you want to hedge. Then there's still a possibility for saving money. The possibility comes from the fact that the reinsurer/over-the-counter dealers also have other pools of businesses. From their perspective, selling you that part of the risk may be actually getting rid of some of the risks on their books that would have been incurred from somebody else. Therefore, they're prepared to let it go at a price that is reasonable. As such, having a view helps increase these savings.

MR. DANIEL R. PATTERSON: Marshall, how were you modeling basis risk? Maybe you can give me some help. Looking at historical data, it's not uncommon for the S&P to be up six percent a year and one of our funds to be down six percent. That 12 percent is going to reduce a lot of negative distributable earnings, that under a two-RBC model, I think we need to try to capture that. I've wondered, is it going to actually reintroduce more risk with basis risk?

MR: GREENBAUM: Yes, I was able to leverage Constellation's experience with basis risk to simplify the calculations that I modeled. At Constellation, we receive fees for over 800 different funds. They are mostly actively managed. That's at the core of our business, the management of basis risk. Because we have such a large diversified group, we're able to manage the basis risk down to the two percent residual volatility level. We have a historical track record after applying out a delta hedging program where we consistently look to optimize the betas and the coefficients against the indices that we're using.

It's a continually optimized type of delta hedging program that we're able to get down to two percent. There are some interesting reasons why that even with a large enough portfolio, from a theoretical basis and under CAPM theory you should be able to get the residual volatility down to zero percent. From a practical standpoint, it's virtually impossible. There are a number of systematic effects that are happening in the active managed fund industry that seem to have all the funds acting and behaving on the same type of information that create a level of risk that can not be diversified away.

When we look at individual fund complexes and over 30 different clients, we are able to measure the residual volatility at the complex level and they experience from 3-10 percent volatility. For the most part, what we've observed is pretty much a normal distribution.

I tried to take an average client that would be representative of funds that would be a typical mix in a variable annuity product. I just earmarked the residual volatility at five percent and assume it was normally distributed around the market value of the GMBD option that the program hedged against. I did nothing more than that. You can extend the model to make the distribution something other than normal to perhaps include the market gapping risk.

MR. MATTHEW COLEMAN: I do not fully understand the quasi-static replication concept. I wonder if you could comment on buying over–the-counter options or getting involved in reinsurance for a small to medium EIA writer. In my view, those are some of the only options for risk hedging. You've offered a second or a third option. At what point of critical mass does static replication or this method you've suggested become a benefit?

MR: RAVINDRAN: In terms of critical mass, in terms of the size of business, it's exactly the same as the over-the-counter contracts. For the savings I showed you, 35 basis points, this is no different from the other sizes that a hedger would typically execute (i.e. there is no impact on the size of the trades).

MR. COLEMAN: But there certainly would be a cost of getting up to speed on this new program, which would require a certain critical mass.

MR. RAVINDRAN: If you have the right template and the right process, one person is enough. Once you have the template and the process, it's very easy and becomes a no-brainer. As such, the plan for anyone interested would be to set up the road map so that it doesn't take more than one resource to do that.

MR. COLEMAN: Is this something that a $100 million a year writer would be able to efficiently use, or would that be $500 billion or what?

MR. RAVINDRAN: The template can be used for anything. The size doesn't matter. But the bigger your portfolio, the more dollar value the savings translates into. Thirty-five basis points on $100 million in savings is a lot smaller than 35 basis points on a billion.