RECORD, Volume 26, No. 3*

Chicago Annual Meeting
October 15-18, 2000

Session 7PD
Hot Topics in Separate Account Products

Track: Investment

Moderator: PETER D. TILLEY
Panelists: MARSHALL C. GREENBAUM
MARY ROSALYN HARDY
HANS U. GERBER
Recorder: PETER D. TILLEY

Summary: The panel discusses current events in the separate account product arena. Topics include:

- *Quantifying the cost of guarantees such as minimum death benefits and minimum living benefits*
- *Equity-indexed product investing and hedging*

Mr. Peter D. Tilley: We have a distinguished panel and Marshall Greenbaum is our first speaker. Marshall is a chartered financial analyst and a certified financial risk manager. He's with Ernst & Young's Actuarial Services Group specializing in financial and actuarial risk management within the insurance industry. He consults with life insurers with risk management issues predominantly utilizing capital market-based solutions. He works extensively with asset/liability modeling and has developed economic simulation models for numerous clients. More recently, his work has focused on assisting clients with derivatives hedging strategies to mitigate market risk embedded within their products. He'll be speaking this morning on hot topics in insurance products, including hedging strategies.

Our next speaker will be Dr. Hans Gerber from the University of Lausanne in Switzerland. He received his doctorate in mathematics from the Swiss Federal Institute of Technology. He was a professor at University of Michigan in '81 before he moved to the University of Lausanne in Switzerland, were he's worked for almost 20 years. I downloaded a list from the Web of Dr. Gerber's publications, thinking I would do a quick run-through, but it ran eight pages so I will just hit the high points. He's a co-author of a text, *Actuarial Mathematics*, with Bower, Beekman, Jones, and Nesbitt. He's an editor or associate editor of several journals including the *North American Actuarial Journal*, and he has won a number of prizes, including the Holmsted Prize, and the SOA's Annual Prize, both those prizes with Elias Shiu. Dr. Gerber is going to be speaking on the paper he co-authored with one of his students, Gerard Pafumi, which was published in the April 2001 issue of the *North American Actuarial Journal*, Volume 4. The title of the paper is, "Pricing

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

Dynamic Investment Fund Protection," and he will indicate how it can be used in practice.

Our final speaker is Dr. Mary Hardy from University of Waterloo. Dr. Hardy has a degree in pure mathematics from the University of London. She worked in the Government Actuaries Department in London, subsequent to that she went to Heriot Watt University in Edinburgh where she received her Ph.D. In 1997, much to our benefit, she came to the University of Waterloo, where she's an associate professor and associate chair for actuarial science. She's a Fellow of the Institute of Actuaries and an ASA. She'll be speaking on the regime-switching lognormal model.

Mr. Marshall C. Greenbaum: I'm going to talk about equity-indexed annuity (EIA) products. As Peter has mentioned, the title of this session "Hot Topics" is appropriate because it gives me the opportunity to talk about some recent activities within insurance products, specifically as they relate to the investment field, but it also gives me a lot of leeway.

I've focused my presentation on equity market risk, which is something that the insurance industry has been confronted with over the last several years with the advent of EIA products, as well as the tremendous success of variable annuities (VAs).

One product I'm going to talk about are VAs and their tremendous sales success. It's interesting that a lot of companies that I've worked with over the last couple of years that had a VA product haven't paid much attention to it. With the tremendous sales success over the last couple of years, suddenly companies have refocused their energy on their VA product. I think there are a couple of reasons for that. One is, it is not just hot guaranteed equity features that are being offered, but what I refer to as sizzling guaranteed benefit features that provide policyholders with extremely valuable guarantees.

What are these features? It started off with guaranteed minimum death benefits (GMDBs). Initially, these contracts had a return of premium guarantee which specified upon death that a policyholder would receive the greater of the premium that was put into the policy or the account value at that time. And now what we have is what we call enhanced death benefit guarantees, 5-6% rollups. I've even seen a 7% rollup out there. These take the premium and roll it up at that 6-7% interest, and then there's a benchmark for comparison against the current account value upon death. And just stop right there and think about it, a 6% or 7% rollup, when you have a balanced fund which is 50% equity and 50% balance, the expected return isn't really that much in excess of it. By the time you subtract out the fees you have a pretty valuable benefit.

But we've moved on from guaranteed death benefits. We've moved into something called guaranteed living benefits, even a bigger beast. Guaranteed minimum income benefits (GMIB) are now a hot topic. A lot of companies are looking to roll out GMIB. What this feature does is it allows the policyholder to annuitize their account value at a kind of guaranteed amount, the same as the death benefit. You could have the 5% guaranteed account value applied that you could then annuitize at guaranteed rates. You don't need to die to receive this benefit. You can elect this. It's your option. There are also guaranteed minimum accumulation benefits. If you basically stay in your policy over a specified period of time, the insurer basically guarantees that your account value will maintain some minimum threshold level. It's typically 100% of premium, and a typical deferral period will be something like 10 years.

And we haven't stopped there. Guaranteed minimum withdrawal benefits allow policyholders to withdraw a certain percentage of their account value over a certain period regardless of whether you run out of funds in your account value. And there's one other, which I'm not going to talk about a lot, but I think it's going to be one of the next big, hot items. It is a guarantee that's attached to a variable immediate annuity called guaranteed payout annuity form. This specifies that once you're in the payout mode of a VA that your annual payment will not be below a certain level, typically 85% of the starting premium.

The next product, although I don't know that it's such a hot product, is EIA. The sales figures are sort of dwarfed in comparison to VAs, something like 1/20th of that. I don't know why this is the case. Because when I think of an EIA and how it's really billed to be upside equity market performance when the markets are doing well, as well as the downside protection when the markets aren't doing well; I really don't understand why this isn't a hotter product. I can conjecture that policyholders like the idea of a VA that they can invest in the same mutual funds that they would otherwise. They're relying in a way on the marketing potential that mutual funds have had over the last few years.

Another product, multi-strategy annuities, is really a supermarket product. It's a combination of EIAs and VAs. This allows the policyholder to select different strategies, each strategy being made up of a particular interest-crediting mechanism as well as a particular index. I've seen products that will allow the policyholders to allocate their money into 25 different strategies, different annual ratchet designs, different point-to-point designs, as well as some past designs, where the insurer can segment a certain portion of their general account assets. This allows the policyholders to participate in those returns of their equity. It has a VA feature. It allows the insurer to capture their mortality and expense (M&E) in a similar manner to the way it's done in a VA, sort of a percentage of assets. So, if the general account is doing well, the insurer collects a higher amount based on higher policyholder accounts.

Another nice feature of this multi-strategy annuity is that the guarantee is typically specified at the contract level. So, it allows for diversification among that guarantee which typically isn't seen in pure EIA product where you have one particular strategy that the policyholder is 100% dedicated to.

VA sales have been good over the last few years. Between 1998 and '99 there is over 20% growth. We have over $100 billion dollars in annual sales, and we have a projected figure for 2000, which is based on doubling the first half-year of sales to be $146 billion. It's going to be interesting to see if the recent market turmoil is going to bring that a little bit lower because as I said, it's based on a pure doubling of first half sales. Again, though, when I mention EIAs the latest figure I have is about $6 billion for 1999. For variable immediate annuities I was able to find a figure of roughly $6 billion. I couldn't get any really solid information on multi-strategy annuities, but I think it's fairly safe to say it's in the single digit billions as well.

What are some of the hot topics regarding VAs? I posed them as questions and I think they follow a logical order. There are a number of direct writers out there that perhaps see their sales lagging. Companies that were once in the top 20 in terms of sales production are no longer there and wondering how to get back up there. There are some that were never there that are wondering how to get into the top 20. The first topic, is how do I improve VA sales? I don't know that it's the entire answer, but one answer has been offering these exotic guaranteed benefits, which has been this really hot topic over the last couple of years.

Once you've gone through that process, and you've decided to offer a new, exotic, guaranteed benefit, the next logical topic is, how do I really quantify or price that specific ancillary guaranteed benefit? In the past reinsurance for these benefits was readily available, so there wasn't much analysis that needed to get done. The direct writers, for the most part, passed the reinsurance contract to the policyholder. That's no longer the case. There's been a dry-up in the reinsurance market, and there have been increased prices that have made companies take a look at this benefit and do the analysis themselves.

I've decided I want to offer the benefit. I've quantified the price. Now I'm going to be collecting a price that's usually expressed in basis points and taken out as a percentage of account value annually. The next logical question is, How do I assess the additional risk? I know this is going to introduce a tremendous amount of risk. There may be a baseline scenario where I don't have any expected claims kicked up. Certainly in the dire economic scenarios we're going to have a lot of additional risk, a lot of additional claims being kicked up. And clearly the way to approach it, the way I've approached it, is to do a fully dynamic stochastic Monte Carlo model. This approach lets you take a good look at the full risk profile. When you do that process you need to incorporate everything that's dynamic in the real world into your model, dynamic lapses. All your capital markets must be dynamic. You must have dynamic lapse behavior. And if you go as far as introducing new sales into your projection, you might want to have that be dynamic. Most people feel that sales are going to increase when the market's doing well but are perhaps going to flag significantly as the market takes a downturn.

You may have assessed the risk, you may have decided where it is; then the last question is, What can you really do about it? Reinsurance is certainly one option, but the other option that most companies are looking towards is internal risk-management solutions.

I want to talk in a little bit more detail about the pricing of this guaranteed benefit, and there are a lot of different views, thoughts, and people in different camps about this topic. I want to give you a little bit of the background. There really have been two competing methodologies that I've seen used in practice, one that I refer to as the actuarial methodology and one that is more the capital market-based option pricing technique. They are really the same methodology though because they have the same objective function; that is, they set the price so that the guaranteed revenue will offset the guaranteed claims on an expected present value basis. The way they get there is the assumption set is different. The actuarial methodology, which is the way actuaries have typically approached these problems, is to set the model off on an expectation basis, and the example that I'm referring to is the GMDB with that 6% rollup.

We'll set the capital markets off in an expected path, as well as volatility, discount at perhaps the cost to hurdle rate for that company. And there are two approaches. One is that you would do your stochastic Monte Carlo model, everything on a stochastic basis, but then perhaps establish a price at a certain percentile. You price for that 95% of your scenarios. Your revenue exceeds your claims for that particular benefit and perhaps staying at the 99% level. The other approach may be to have all of your assumptions, your expected return and volatility to be a little bit conservative in nature, and then you can price for more of a mean result. Or you can do a combination of both, conservatism in assumptions as well as conservatism in that percentile that you choose to establish your price.

Now let's just benchmark that or compare that to this other school of thought, which is option-pricing techniques that totally rely upon the Black-Scholes market framework that these recent wise Nobel Prize winners have laid forth for us, which says that you can actually project based on no-arbitrage principles. It says that you could project your funds at the risk-free rate and discount at your risk-free rate to establish sort of a fair market value price for your particular benefit design, and if you do that, you come up with a completely different answer than the actuarial approach. A lot of discussions at other Society meetings have been somewhat dedicated to this. It really would be a mere coincidence that these two approaches would be equal. The assumption set and the thought processes are completely different.

I'm going to talk about an approach that I think makes sense, and I should state that this is my viewpoint and not necessarily the viewpoint of Ernst & Young. I've used a lot of the recent actuarial literature out there, as well as the financial economic theory that we've learned over the past four years, and tried to come up with a methodology that in a way bridges the gap and puts all the competing elements in there. We start off with this notion that the value of a block of business is equal to the average present value of after-tax statutory distributable earnings over a stochastic scenario set. I clearly don't have the time to explain in very much detail why this is the case. Everybody accepts this as a true statement, but David Becker wrote an article which explains essentially that after-tax statutory distributable earnings are the free cash flows to a shareholder. They represent the return to the firm's shareholders and the appropriate risk reward relationship should be embedded in that.

The next critical item we look at is the contract, and we need to recognize that we need to price these higher contracts with and without this embedded guarantee. One of the problems I see with that other methodology is we kind of strip out the guarantee. We just look at the cash flow of that guarantee by itself, and it doesn't reflect any of the interactions that that guarantee has with the base-underlying contract. We haven't sold to a policyholder a stand-alone embedded guarantee. We've sold them a VA contract. What we need to do is take a breath and price the entire contract, but the sum of the pieces isn't going to be equal.

Third, the assumption set that you use in the stochastic model, as well as the discount rate that you use for discounting, is critical. And this is something that has been discussed in detail in a recent article that was published in the *North American Actuarial Journal* by Luke Girard. The title is "The Market Value of Insurance Liabilities: Reconciling the Actual Appraisal and Option Pricing Method." He showed in that paper that if you didn't use an arbitrage-free scenario set that you'd get these nonsensical results. Typically if you do any type of pricing work, and you invest your surplus into a higher expected return asset, that typically tends to drive down your price and increases the risk, and that's sort of the nonsensical answer that is illustrated in his article.

The last item is the discounting rate for the embedded option. What option pricing theory and Black-Scholes learned and established is, if you take a look at your risk-free rates, what you need to do typically is add an option-adjusted spread (OAS) to calculate your price of a product, whether it's a fixed income security or an embedded option. What I do is actually look toward other asset classes for a benchmark to establish that price. That's essentially the one unknown in this pricing framework, that is, What's that required OAS that an investor would need to earn to enter into this line of business? Unfortunately, we have no liquid traded market for that and that's the one subjective piece that we need to take some good educated guesses on.

I'll go through a little bit of a case pricing study that I put together. The way I approach this is, instead of trying to price directly for what the guaranteed benefit should be, I ask, What's currently out there? What's the charge out there? It doesn't make sense to actually go through this pricing exercise and get 80 or 90 basis points. It's irrelevant because if I charge that price in the marketplace, the product's just simply not going to sell. I'm going to do it with and without a GMIB, which is one of these exotic guarantees I've been talking about. I'm going to go through the one without the guarantee first, which is a typical base M&E charge; that's a charge to VAs, roughly 130 basis points, and for this example we're just going to have a plain vanilla return on premium design for the guaranteed death benefit.

I've gone through this theoretical exercise to calculate what's the OAS that I need to establish for the price or the current price that I could basically sell the product at. And I've come up with 10%. Of course, I haven't given you any of the detail or assumptions, and obviously the answer is vastly dependent on the assumption set. The objective is the same as a bond trader or an option trader comparing the OAS to determine if he wants to sell a rich security or buy a cheap security. You would compare the OAS of this product to what he thinks the required OAS would be for an investor in this particular line of business. This is the subjective piece that I was talking about. I've come up with a 3% required OAS. If the product price is in excess of 3%, the product line would be a go, and that's typically how you would look at it, always price the product for an excess OAS over the required OAS required by the investors or the shareholders.

Where did I come up with this 3%? Again, we don't have this liquid actively-traded market, but we have these other fixed-income securities that we can use as benchmarks. I'm looking in this case towards the collateralized mortgage obligation as well as the mortgage-backed security asset classes for benchmarks. And when you do this same exact methodology for those, you tend to solve for reasonably high OASs, from 150 to 250 basis points. The question you might ask yourself is, "Why is there even an OAS?" The reason why that OAS is established is because when people value mortgage-backed securities, they have to use a prepayment model. They can't hedge that risk away, therefore, they have to take an educated guess, and they want to be compensated by a higher required return.

If I try to compare the embedded risk in my VA, which is mortality, although I can't logically guess this, I would have about the same kind of model as in this specification for a prepayment model. And that's where I targeted my 3%. If you actually calculate your after-tax statutory distributable earnings using a 3% return on OAS, you're going to get a value in excess of zero, in this case, $8,700 for $1 million in sales, which is essentially a very good return for this product.

Now let's compare it to the VA with the GMIB. In this case I've thrown in the GMIB, and a reasonable sales target for this is 35 basis points. I'm going to plug that in as what I think I can sell my product at, go through the same exercise and come up with my priced OAS, and in this case it's 9%. What we see is it's a little bit less than the VA without the GMIB, which means that perhaps the cost is in excess of 35 basis points. I'm going to be subsidizing a little bit with the base end of the contract. What might the required OAS be for this product? Maybe it's not 3% anymore because I have an additional model in the specification there that I have to worry about which is this annuitization election.

When you're analyzing GMIB, the big mystery is, "Are people going to elect this benefit?" Their account value is going to be down. Over 10 years it's going to be 20-30% in the money. I have to take a good stab at what percentage of the population remaining at the time is going to elect. I need to be compensated for that with an increased required OAS. When I go up there and calculate my after-tax statutory distributable earnings as my risk-free rate plus 5%, I get a $5,200 value for $1,000,000 in sales, which is less than the VA without GMIB, which makes logical sense. You're offering a guaranteed benefit that's going to eat into your base M&E a little bit.

We can't stop there because that's not the entire picture. We need to come up with a projection of how much volume these products will generate. And the product with the GMIB is going to be so much more saleable because it has that exotic guarantee that policyholders might be looking for. And I'm going to be able to sell five times the amount of the plain vanilla VA. The last wrinkle is even though my return over my required OAS is only 4% (versus the VA with the GMIB at 7%), I could sell five times the value. It's going to create close to four times the fair market value.

The next question is, What do you do about the risk? The short answer is that you could hedge with reinsurance. You simply can take out the claims risk if there's available reinsurance or you can establish this internal risk management-hedging program that would use financial derivatives such as indexed futures and option contracts. You would establish a portfolio that would give you the exact offsetting impact that the claims that you've given to the policyholder creates.

As far as the GMDB and GMIBs, when the market goes down you're going to be hit with claims, but perhaps a direct writer would go out and buy S&P index put options, which would give him the offsetting position. That's basically the short answer. Obviously, there's an infinite number of different hedging strategies, approaches, and ways to get about that, using static or dynamic hedging. One thing I wanted to note was that even without GMDBs and these exotic guarantees that I've been talking about, there's incredible variation in your risk profile curve which is just based on your mortality and expense risk. The way a VA works is that an insurer will receive more money when the market goes up on a dollar basis and less when the market goes down. There actually can be a loss in those down scenarios when your revenue gets bumped up against your fixed expenses. It was once believed that VAs were a riskless product, and now everybody's found out that's really not the case.

How do you assess the risk? I think one of the best ways to look at it is to put a complete risk profile curve together. You would employ a stochastic Monte Carlo model, and over that you would basically take a look at some sort of summary statistic, here it's PVD, and take a look at the different percentiles. This is the theoretical market value that I established. We're working with the VA with the GMIB that was $5.2 million for $1 billion in sales. That would be a theoretical fair market value that would securitize, push it off to the capital market, if there was a liquid traded market for that particular risk, but we don't have that.

You have a couple of other choices. You go naked which, in the financial world, means you don't hedge the put. You sell the product, you charge the price that you determine is appropriate, and then you sit with it. It is quite variable, and there's a large tail where there are a lot of outcomes where you don't earn in excess of your hurdle rate. Then there's this other approach, which is dynamic hedging. In a perfect world you would be able to dynamically hedge for this theoretical market value, but in the real world we have a bunch of real-world intricacies and complications that don't allow us to achieve that.

We have things like basis risk, which means we're going to have to enter into derivative contracts, that don't match the subaccounts that are embedded in a VA. We have this election risk. We have to take a good stab at how many people are going to elect the GMIB. Could be off. Could be right. We have mortality risk. And one last thing is that in the real world you just can't rebalance continuously, which is how the theory is set up. There's going to be some residual risk just because in the real world we're going to implement the program. We're going to need to rebalance daily, weekly, monthly, if we're lucky.

I work with a number of clients that are hedging their EIA. I hate to say it, but I've really failed to come up with one that I think is doing it in an optimal manner. Certainly there are a number that are doing very good jobs and very acceptable jobs, but there's one pretty big issue about hedging EIA products that I want to bring up. I almost want to hypnotize the audience a little bit. I'm going to use an expression: "borrow the cash." You need to borrow the cash when you hedge your EIA product, and, unfortunately, there are a number of writers out there that aren't doing that. They're hedging what I call the credited interest mechanism, which is this fictional, notional amount that they use to bump up the policyholder's account value, which doesn't necessarily turn into cash. You know the policy is going to turn into cash when that person exits the policy.

The practice I've seen is that with the annual ratchet product, they invest 100%. Everybody is there in the projection. They don't reflect anybody who's going to lapse. If anybody lapses, and they don't have to credit that person any kind of interest mechanism there, they sort of overhedge. They bought a call option that they didn't need and they're basically taking a bath and that basically creates equity market exposure.

I'm going to go through what I think is clearly one of the most common designs from an annual ratchet product. In the annual ratchet product an insurer will credit the policyholder a percentage of an equity market index, which is typically the S&P 500, and that percentage is usually referred to as the participation rate (PR). An insurer's profit isn't always in the PR. Sometimes they offer a spread where they subtract a value, and sometimes they enter these caps on a return as a way to get profit. I'm going to talk a little bit about this PR option in which all insurers now retain the right to adjust this PR on an annual basis so that if there's a change in market conditions, interest rates move up or down, they can change the PR in their product. This has, I think, led many insurers to the wrong conclusion that they don't have to worry about the options that they essentially sold in the policy after that first year. They concentrate on hedging that first-year interest mechanism and let the other stuff float, and that's clearly, in my mind, the wrong way to approach this.

Assume that you have a five-year annual ratchet product, and you were going to go out and buy this one-year option that is worth five cents. Assume that's a 40% PR, a 6% discount rate. The guarantee is 90% of premium at 3%, and the base lapses are assumed to be 2%. Mortality is 1% a year. And everything else is constant. It turns out in this case if you mark to market those forward options at five cents, and you have two, three, four, and five. There's a fair market value established, it's not zero. You've actually given the policyholder a deferred one-year option that you need to account for in my mind. In this case the total option that you have to purchase for a policyholder is 21 cents. In year one it's only five cents, which leaves a balance of 16.1 cents for the remaining options you're going to have to purchase for the policyholder.

Let's go through a particular scenario where the equity-indexed market goes up 10%, and, as a consequence, you feel that the markets are doing well. You're not going to get expected lapses. Lapses go from 2% to 1%. Or, if I were discounting and discounting into the option cost, and we go through the same exact calculation, the market goes up 10%, so that first-year option becomes extremely more valuable at five cents to 7.6, but there's also a change in the market value as options change. More people will be around to collect on that option. The year two-to-five option now becomes 16.3 cents, not 16.1 cents anymore. And in my mind you should hedge that because you can have substantial changes in those option costs in year two through five as the fair market value for the statistics that go into that calculation change (such as implied volatility).

Another example of where the hedging doesn't work out is when you do the one-year mechanism. Assume that the market has three years of negative returns. After three years of negative returns in the policyholder's account, you still have $1, and the guarantee starts to kick in, which means you no longer need to buy the option worth five cents because the strike is going to be greater than one. It's going to be less than five cents. And these are all the elements I want to talk about: borrow the cash. You need to properly reflect your option sensitivity when you go out and hedge this. The natural result is that it leads to something called dynamic hedging. You really want to appropriately reflect that because there won't be actively-traded options to exactly offset this design.

You need to calculate the option contract on the full horizon. You can't take a look at pieces because we can't take out pieces. If I wanted you to bring home one issue from my presentation, it is that you need to take this macro approach and value the entire contract. One last comment: In my mind this is the way that you should bifurcate the liability for FAS 133. Take out all those options. I think it's completely ludicrous to mark to market one-year credited options for the first year. If you look at an option where you have one date of expiration and you completely ignore the other option, you're not showing the whole picture.

Mr. Paul Hickman: Marshall, you mentioned your opinion that a FAS 133 treatment of the EIA should require a pricing of the renewal options. In your opinion should that be based on the underlying policy guarantee, maybe 20-30%, or should it be based on what you feel that you're going to need to pay to keep the product competitive in the market?

Mr. Greenbaum: You're asking about how you go about adjusting the participation rate in those forward-looking options?

Mr. Paul Hickman: Yes. A lot of policies have like a minimal guarantee, anywhere from 10-30%, but, probably your current sales certainly will have to be higher than that.

Mr. Greenbaum: I think you reasonably have to adjust your participation rate. There are some contracts where you have the right to set a PR of zero. If you recognize that rate, and you recognize that obligation in those deferred options, then clearly you have a cause for determining that your option costs in year two through five are zero. Now, the reality of the situation is that it's going to be difficult to implement such a strategy even if you intended to implement that strategy. In my view, you need to realistically reflect your dynamic participation rate strategy in the calculation of those forward-looking options. And you need to reasonably reflect economic pressures, so you may want to actually levelize your option cost as a percentage of account value and exactly offset, but you're not going to be able to do that because of competitive reasons. I think you should even go as far as reflecting that. I think what should be reflected is your best intent for going forward and what you would reasonably do under those particular scenarios.

From the Floor: Marshall, you had a 35 basis point charted for the GMDB. And you had a differential with and without the GMDB.

Mr. Greenbaum: In the example I discussed, it would be more than 35 basis points to get back to the same priced OAS at 10%. I actually didn't go through that calculation, but it vastly depends on your assumptions for lapse behavior as well as the election, so there's actually a wide range that the result can reside in, from as low as 25-30 basis points to close to triple digits.

From the Floor: Additional?

Mr. Greenbaum: Additional, as the rider charge. In some of the work, I use what I feel are reasonable lapse expectations as well as my expectation for annuitizations. This is based off a formula that's based off of how much the individual is in the money; it tends to reside between 30 and 50 basis points.

Mr. Hans Gerber: I want to give you some basic ideas on a paper I co-authored titled "Pricing Dynamic Investment Fund Protection" that was published in the April 2000 issue of the *North American Actuarial Journal*.

If you were to look at the fund, and F(t) is my notation for the initial fund unit value at time T. Typically, fund unit value fluctuates in time. This is a typical possibility, called a trajectory. Normally you think it's a stochastic process. Now, for the basic ideas that's all we need, but to get some concrete results we have to make some assumptions, and, in fact, we will make the usual assumption that F(t) is a geometric Brownian motion like in the Black-Scholes model. Whenever I say something specific, numbers and so on, it will be based on this assumption, but I guess the idea is not to depend on this.

We can distinguish two cases. One would be what I call a static protection, and the other is what I would call a dynamic protection. The simplest way to obtain the static protection was simply to buy a European put option with say an exercise, a strike price, K, and for the date in which we're interested, T (for example, one year from now). Well, then instead of having F(T), at maturity we will have the bigger of F(T) or K. So, K would be a guaranteed value. Simple enough. Of course, we have to take into account the cost of it. Subtract the cost.

Here is the drawback of this static protection. In certain situations, if you have bad luck, you will never have more than K at the end. K would be guaranteed, and if we have a put option, yes, you will have K, but there are no illusions.

Now with the dynamic protection that I will present, at any point in time you will have more than K. The outlook is more optimistic, but, of course, the cost will be higher. Let me get into the dynamic protection. The idea is to replace the original fund unit value that is F(t), by what we call the upgraded fund value $\tilde{F}$ (t). This is not just for at maturity but at all time, small t, before maturity. F (t) would be called a naked fund, and this would be called the protected fund or really the unit value.

The whole question is, "How would $\tilde{F}$ be defined?" First, I'm going to show you intuitively and then give the formula. Whenever the unit value falls to level K, and K is like a guaranteed level, you add just enough money so that the unit value does not fall below this level, if you are outside of the rate of return, which has been the same as before.

This would be F(t), the original fund unit value, but when you hit the barrier, now enough money is added. That's the guarantee that you never fall below. You will again get off this. In the terminology of stochastic processes this is like a reflecting barrier. That's the whole idea. What I have here, is actually an $\tilde{F}$(t). After this point they are all the same, but afterwards it's the operated fund unit value.

Figure 1 is a typical sample path of the fund unit values. Initially they start out with the original fund unit value. Afterwards we have F(t), which would be below K sometimes, while $\tilde{F}$ (t) is always above K. This is the formula for the precise definition:

$$\tilde{F}(t) = F(t)\,\max\left\{1, \max_{0\le s\le t} \frac{K}{F(s)}\right\}$$

for $0 < t \le T$

The upgraded fund unit value of $\tilde{F}$ (t) is the original, the naked fund unit value F(t) times a factor greater than or equal to one. It's the maximum between one and the maximum over all values of s of K divided by F(s). In other words it's the ratio of the guaranteed level divided by the observed fund unit value of the time t, and take the maximum. These are actually the mathematical definitions.

Now, in the paper, assuming that F(t) follows a geometric Brownian motion, there's an explicit formula for the cost of this protection. How much does it cost to go from F(t) to $\tilde{F}$ (t)? There's an explicit formula; just a notation. For the price it will be V (f,T) where f is the initial value, and T would be the maturity date. Well, it also depends, of course, on sigma, the volatility on the risk-free interest rate. We don't want to look at the formula. Here's an example of what that means in terms of the

$$\tilde{F}(t) = F(t)\,\max\left\{1, \max_{0\le s\le t} \frac{K e^{g t}}{F(s)}\right\}$$

cost. For example, always initial unit values of 100, volatility, 0.2, risk-free rate of interest, 0.04. And different values of K. So, if K is 80%, it's much cheaper than if K is 100%. In each line this is increasing.

TABLE 1

Price of the Protection with $f = 100,\ s = 0.2\ and\ r = 0.04$

T\K	80	85	90	95	100
1/12	0.0001	0.0065	0.1304	1.0797	4.5189
2/12	0.0109	0.1093	0.6338	2.3761	6.3359
3/12	0.0624	0.3340	1.2463	3.4770	7.7069
4/12	0.1626	0.6313	1.8676	4.4370	8.8463
5/12	0.3035	0.9659	2.4706	5.2943	9.8376
6/12	0.4746	1.3180	3.0481	6.0732	10.7233
1	1.7709	3.4239	6.0120	9.7476	14.7931
2	4.4061	6.9231	10.3118	14.6840	20.1295
5	10.1373	13.7031	18.0257	23.1640	29.1716
10	15.6391	19.8688	24.7909	30.4504	36.8905
20	20.8713	25.5995	30.9834	37.0626	43.8762
	25.6000	30.7063	36.4500	42.8688	50.0000

This was the first possibility. The idea was to have a constant guarantee to level K, but perhaps it's desirable to have a guarantee, which is a function of time. The next function of time, means essentially it's a guaranteed rate of return, gamma. What the guarantee is that the unit value at any time t does not fall below this barrier, which is: $K e^{s t}$

The upgraded unit value at time t is the original unit value times t times this factor, and the factor now is:

$$\tilde{F}(t) = F(t) * \max\left\{1, \max_{0<s<t} \frac{Ke^{g t}}{F(s)}\right\}$$

Now, as it turns out, calculating the cost of this protection is not more difficult than to calculate the protection with a constant barrier. The rule is simply this:

- Cost for protection with $g > 0$

  =

- Cost for protection $g = 0$ with

The cost for the protection if gamma is positive, guaranteed rate of gamma, can use the same formula as before. It's the cost for protection of gamma equals zero, but we have to replace r, the risk-free rate of return just by the difference (r-g). The calculations are easy enough also.

Let me just make two final remarks. One, while there is an authority formulation, if you like an authoritative design, how I presented it, was in terms of F(t), the naked fund unit value, and then the upgraded fund unit value, $\tilde{F}$(t). It's also possible just to talk about one fund unit value, but then as a consequence of the guarantee you increase the number of units.

Alternative Formulation

F (t) = F(t)

$$\tilde{F}(t) = h^{(t)}\ F(t)$$

h (t) number of units

$$h(t) = \max\left\{1, \max_{0 \le s \le t} \frac{K e^{g s}}{F(s)}\right\}$$

$$0 \le t \le T$$

In other words you can introduce h (t) number of units. So, originally you would have one unit, and the value is F (t), and then with the protection we would have h (t). Instead of that you would simply have h (t), instead of t. That's just another way to predict the value. That's the advantage that is on the one fund unit value, but the numbers of fund units would be increased. Well, the formula is simple enough.

Then the second final remark would be that, other designs, such as dynamic designs, are a path-dependent protection (Figure 2). For example, the guaranteed level would be a fraction of the observed maximum. M (t) would be the observed maximum of the time t, and at any time you guarantee a fraction fee of that observed maximum. It is more complicated to value this, and it may cost also. Actually, I have seen a picture like that on the home page of all of the Swiss banks of CS.

Ms. Mary Rosalyn Hardy: My talk is complementary to the other two speakers. I'm talking about the same kind of products, but I am interested in model risk. In both cases, I think that they mention they're using a standard Black-Scholes framework, including assuming a geometric Brownian motion or a log, which is equivalent for the indiscrete time, assuming a lognormal model for stock returns.

In looking at model risk, this lognormal model for separate account products (and I've been involved in the separate account products in Canada) is so wrong. Now, if the lognormal model is the wrong model, and you have a typical option of maybe six months, a long-term option in the options market, would be a year, it's not going to cause you too many problems. But, if your lognormal model is wrong and you have a 10- or 40-year contract, then the differences are going to accumulate, and you're going to be in trouble.

What I've been looking at is applying the regime-switching model to separate account products and Steve Craighead stressed that you might be interested in that. The idea of a regime-switching model was developed by an econometrician, J.D. Hamilton, in several papers. This is a picture of the S&P. What people have been worried about with model risk or model error in modeling something like the S&P with a lognormal model are these clusters of high volatility, and this isn't real volatility. This is a rolling 12-month average volatility. First, we have clusters of high volatility, and, second, they're associated with very poor markets. This is the oil crisis and the '87 crash, of course. Neither of these things is reflected in the lognormal model or the geometric Brownian motion. Earlier ARCH and GARCH models do model clusters of volatility but at a cost of some complexity.

I just thought I'd put in a little warning, that if I model starting in 1956, which is convenient for me because it's when the Canadian index, the TSE, started, I get very different results than if I go all the way back to 1930. I get very different results, again, if I do what many insurers are trying to do which is start in 1975 or 1980 on the grounds that that's recent. My feeling is that anything you do has to at least include the '73 crisis because for these sort of guarantees we're looking at in Canada that would be the most serious scenario. And the economic outlook or the outlook at the moment is that another oil crisis is certainly not impossible, but if you go back to 1930, and also include the Depression, you get depressing numbers.

The independent lognormal model is commonly used, most people will know about it. I compare it with the regime-switching model when you look at the log of the accumulation stock returns, the log of the accumulation factors, and see that these are independent, identically distributed. I'm actually looking at it in months, but geometric Brownian motion. You can divide it into any period. Non-overlapping periods are all independent, and the log returns are normal. In other words, we can model them as a mu plus sigma times some normal unit variate.

And this is a nice model. Financial economists have been relying on it for ages. We can do lots of mathematics with it, and it is very easy to implement analytically or in stochastic simulation. The trouble is that however nice it is, it doesn't actually fit the data, and that's always a drawback. Data's such a pain. In particular, autocorrelation and volatility bunching doesn't fit well. Hamilton looked at the whole school of regime switching as a principle, and I'm just looking at the lognormal version of it, but we basically say there are two modes for the market. It's either in frantic panic mode because things are going badly, and the volatility's very high, and the returns are negative, or it's in a kind of stable, ongoing mode, and returns are higher, and the volatility's much lower. Most of the time it stays in the stable mode, but from time to time it jumps into this unstable mode, and it stays there for a while and then jumps back again after a few months as the period of uncertainty passes.

We use two regimes because, in most cases, that seems to adequately describe the data. There is an argument for three regimes also, but the two regimes are an output of a fitting process, not an input to it. It's jumping the Markovian Way. That means we're assuming that the process, jumping from the high volatility regime to the low volatility regime, and vice versa, depends only on where it is, not on how it got there, nor on how long it's been there.

What this means is that the one that I call Regime 1 has low volatility, higher returns, and high persistence, meaning it tends to stay in that regime a lot longer. Regime 2 has high volatility, low returns, and low persistence. This automatically includes the bunching and the association between high volatility and poor returns. This tells you what regime you're in. It's either equal to one or two. And it just says that our log returns depend on the regime, and they're going to be normal, but the parameters are going to be different for the two regimes. Conditionally these things are independent, but unconditionally we have some auto-correlation because of the association from one month to the next with being in one regime or the other, and we have transition matrix and transition probabilities, jumping between these two regimes.

Why go to the trouble of two regimes? Well, it actually fits the data better. It reproduces that volatility bunching. It catches extreme observations in terms of the month. I've just used it for the monthly data although I looked at daily data as well. For example, the monthly, for the 1987 crash, which is some ridiculous number like 23 standard deviations away from the mean in the lognormal model is still an extreme observation, but there's a 5% probability that we would see a 1987 crash within the '56-99 data set. In other words, the 1987 is valuable. It's extreme. It's by no means out of the range of what we'd expect to see. In 5% of similar length simulations we would expect to see a 1987 type crash. It's quite an intuitive model because we all know that the market suddenly gets jumpy and then sort of relatively suddenly doesn't again, and it's different. For those of you who know the ARCH model, it's different than GARCH where GARCH assumes that you might suddenly get jumpy, but then you gently go back to normal, and the idea that you jump back to normal actually seems to fit the data better. We can get analytical results, which are always nice for academics, but, probably more important, it's really easy to simulate. It's an extra step in simulation from the lognormal because each time you have to simulate which regime you're in, and then you're just simulating a lognormal process.

It's a bit more complicated (a disadvantage). When I updated the data to include the past year, there was some marginal evidence that three regimes might be better for the S&P 500 total data that we can still do everything with, but it's less attractive because we're jumping from a six-parameter model to a 12-parameter model.

We can get maximum likelihood parameters for this process. It's pretty easy with spreadsheets, as my students explained to me after I carefully did it in more complicated ways. These are the maximum likelihood parameters, TSE Canadian index and the S&P 500.

They're pretty similar, and the interesting things to notice here are the errors. These are estimated standard errors, and the errors for the Regime 2 figures are much bigger because we see a lot less of Regime 2. We're only in Regime 2 for about 15% of the time. We're in Regime 1 for about 85% of the time. When we look at the data it gives us less information about Regime 2. In particular, these transition probabilities have high standard areas associated with them.

Figure 3 shows you what the effect is over 10 years. If you look at the accumulated proceeds of a 10-year unit investment with the TSE parameters, the thin line, this is the density function for that. The probability of doing badly is much greater for this lognormal line than it is for the thinner. It's much greater for the regime-switching line than it is for the thinner lognormal line. We've increased the fatness on this side, and we've also a little bit increased the fatness of the tail on that side. We could do worse; we could do better.

One of the reasons this seems to be a nice model to use is that it does actually seem to fit better than most other models that are in common use. Model selection is a sort of combination of art and science. Looking at some of the selection criteria, what do we look at? One of the things we look at is the log likelihood. The likelihood is proportional to the probability that we would have seen in this model, given the parameters.

The number of parameters is an issue. In general we'd rather have a simpler model than a more complicated model. If we're going to add parameters, they have to pay for themselves by really improving the fit substantially. Three common selection criteria are the likelihood ratio tests. It tests whether the extra parameters that you add in a model pay for themselves. And the Akaike Information Criterion and the Schwartz-Bayes Criterion, commonly used, have very simple criteria for selecting a model. We look at a function of the log likelihood and a number of parameters. I throw that in just to give you the message that for the TSE this is building up by number of parameters.

Table 2
Maximum Likelihood Parameters, with Estimated Standard Errors

TSE 300		
$\hat{m}_1$ = 0.012(0.002	$\hat{s}_1$ =0.035(0.001)	$\hat{r}_{1,2}$ =0.037(0.012)
$\hat{m}_2$ = -0.016(0.010)	$\hat{s}_2$ =0.078(0.009)	$\hat{r}_{2,1}$ =0.210(0.086)
S & P 500		
$\hat{m}_1$ = 0.013(0.002)	$\hat{s}_1$ =0.035(0.001)	$\hat{r}_{1,2}$ =0.040(0.015)
$\hat{m}_2$ =-0.18(0.014)	$\hat{s}_2$ =0.075(0.009)	$\hat{r}_{2,1}$ =0.380(0.123)

What Table 2 tells you is, that compared with the independent lognormal, autoregressive order one, ARCH, GARCH, a regime-switching autoregressive model, and a regime switching lognormal with three regimes, with the two regime, regime-switching lognormal actually does best. It does best in Schwartz-Bayes, Akaike, and in the likelihood ratio test for the models above it, fewer parameters. H(0), the null hypothesis, we use them as models. The more parameters, the null hypothesis, we use them, and in all cases it seems to fit better. One of my colleagues in the task force in Canada is very keen on stable distributions, which is really horrible to use, but it also fits better than stable distributions, and we don't have to use them.

There's marginal evidence here that for the S&P the regime-switching does better in Schwartz with two regimes, but we need an extra regime, will do a bit better on Akaike and the marginal evidence on the likelihood ratio test. And the third regime is an ultra stable regime. It's an ultra low volatility. But I haven't used it. I decided that Schwartz knew what he was doing, and I didn't like using the extra six parameters, but nothing that I do can't be done with three regimes.

I said we can do things analytically. We always like that in academia, because, if you fix a period (say 10 years, 120 months), once you condition on the total number of months you spent in Regime 1, say, so all the rest of the time was spent in Regime 2, you just have a lognormal model. This is a beautiful thing about the normal distribution, if you have normal distribution, you get another normal distribution. In our log returns we're just adding R – lots of Regime 1 to n minus R lots of Regime 2.

Conditionally we have a lognormal model and that means if I know the distribution of R, which is a random variable that depends on this Markov jumping probability, or the two; one that says you go from one to two, and another one that says you go from two to one, if I know that distribution, I can do anything with this distribution. I can get any probabilities out. I can get any moments out. I can get Black-Scholes prices out. Once I derive this probability, which is reasonably straightforward, I can get all these other things. Anything I can do analytically with a lognormal distribution I can also do analytically with this distribution, which is a nice result.

Applications

We're here to talk about guarantees, separate account guarantees, and we're very interested in those in Canada. Let's look at a simple guarantee, 100% of premium on either death or maturity. This is a relatively common Canadian contract term. You have an *n* month contract. You're guaranteed 100% of premium if it matures through death or getting to the end of the term. The underlying stock price is in this regime-switching lognormal model, and then we have a fund, which is after deduction of the charge for management expense ratio.

As Marshall mentioned, the kind of actuarial approach to this is to assume a passive investment strategy. Stick a chunk of assets in bonds and hold them with the idea that in 95% of cases, or thereabouts, they'll be enough to pay any guarantee that applies. And we've been looking at that approach, and that approach is by far the most common we're doing with these contracts in Canada. And there's also, as Marshall mentioned, the hedging approach where you calculate a hedge portfolio and project the costs, and you have to do that on a risk-neutral basis, which you can do either on a straight lognormal or on a regime switching lognormal.

I've calculated the hedge portfolio using straight Black-Scholes, a single regime lognormal model. Then I projected in real-world measure the hedging costs under a regime-switching lognormal, so that model error is going to come out in these extra hedging costs in the P-world, in the P measure simulation, in the real-world simulation. We simulate the additional cost from the hedging that happens monthly from the difference in the stock return model and from the transaction costs.

I see Marshall's a bit worried about monthly hedging. We could do this more frequently, but, in fact, don't. We can, for example, simulate and get a quantile reserve. The idea is that in both of these cases either use the actuarial approach or use the hedging approach. You need a good, real-world measure of stocks to identify your costs, either in the actuarial approach your total costs and in the hedging approach your costs over and above the hedge that you maintain to meet the guarantee.

Guaranteed minimum maturity benefit (GMMB) is the payoff when you mature. And we just calculate the present value of future loss. This is a random variable. Mortality is deterministic, which works all right. It is that you survive to get your payoff, which is the maximum of the fund and the guarantee discounted minus the present value of your expenses coming in. That's the actuarial approach. We have this random variable. We can simulate it and do any calculations we want to determine how to get from the present value of future loss to a price or a capital requirement, and capital requirement's actually what I've been working on.

For example, you could use a quantile, i.e., just take the 95th percentile. What we've actually been doing is the conditional tail expectation (CTE), which instead looking at the 95th percentile simulation, we're going to take the average of the worst 500 outcomes for a 95% conditional tail expectation. Don't worry too much about this. This is the easiest intuitive way to think about it or the more robust measure of the risk in the tail. If I take this right to the end, if I take this conditional tail expectation with a parameter zero, I just get the mean of the expected loss.

We can do the same for the dynamic hedging approach. Again we can develop a present value of the future loss, allowing for mortality, and this can allow for mortality on lapses. You have the hedge. This is the hedge portfolio under Black-Scholes, and these are hedging errors and transaction costs and management charge income, all discounted, because of the result, the error from discreet hedging and also the error because we're using a more realistic model here than we are in H(0).

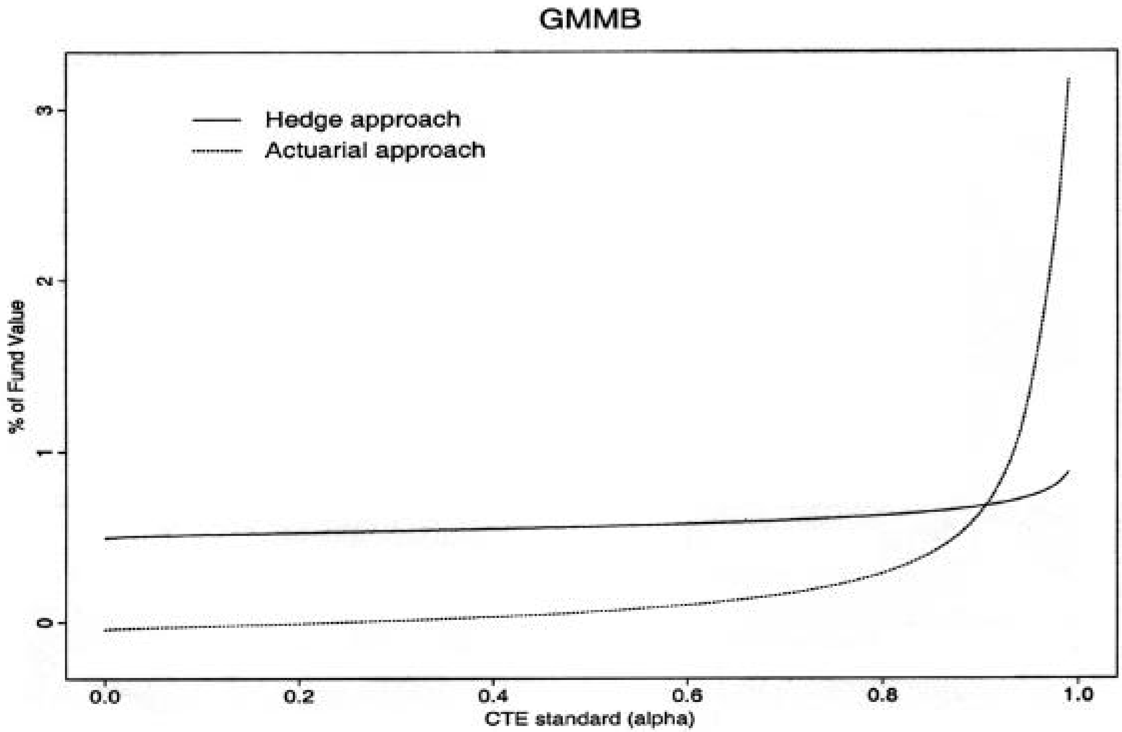

So, in both cases we can get the present value of future loss, and in both cases we can calculate the CTE. I'll just put up the results here just to show you GMMB. This is the CTE. This is our measure of tail risk, which we're going to use as a capital requirement up around here somewhere. This is going to be the capital requirement for these contracts. This is your CTE, your tail risk, if you use the actuarial approach. It's not entirely passive because your probabilities are going to change. You're going to have to change that fund as a result of the experience. What happens with the hedge approach, and at this end we get the mean. This is our mean cost. It's cheaper on average to take the actuarial approach. The hedging approach is more expensive on average. What you pay out you hedge. In 95% of cases you don't need it, and you've wasted your money. The hedge doesn't pay off. It's some cost. Whereas, in this case you stick the money in the bank account, and if you don't need it, you take it back again.

This does not allow for the cost of capital, depending on the capital requirement. This is just looking at the liabilities. At this end the risk at the far end is much greater under the actuarial approach than the hedging approach. And we get the same sort of picture when we look at the GMDB, which is a very similar picture, which is a fixed return of premium GMDB.

Figure 1
A Typical Sample Path of the Fund Unit Values

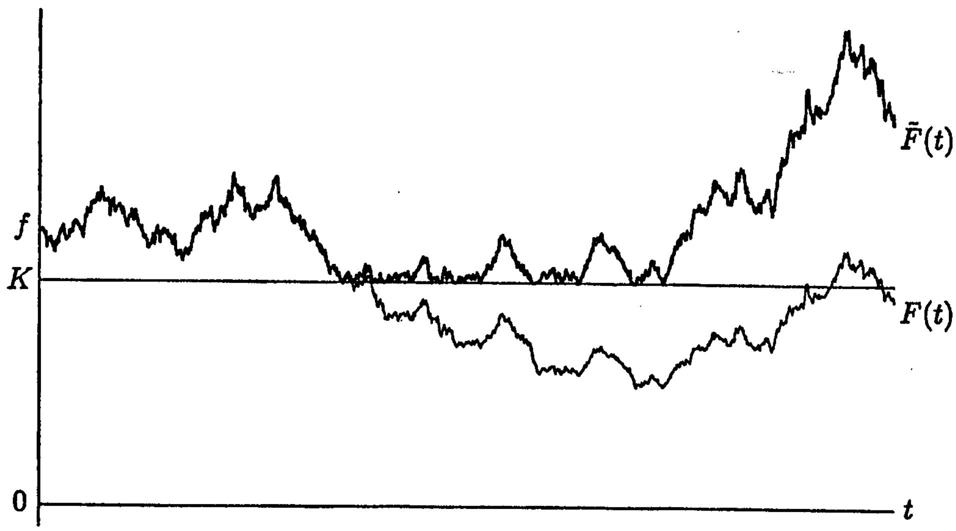

Figure 2
Price of the Protection as a Function of *T* with $f = 100$, $\sigma = 0.2$, and $r = 0.04$

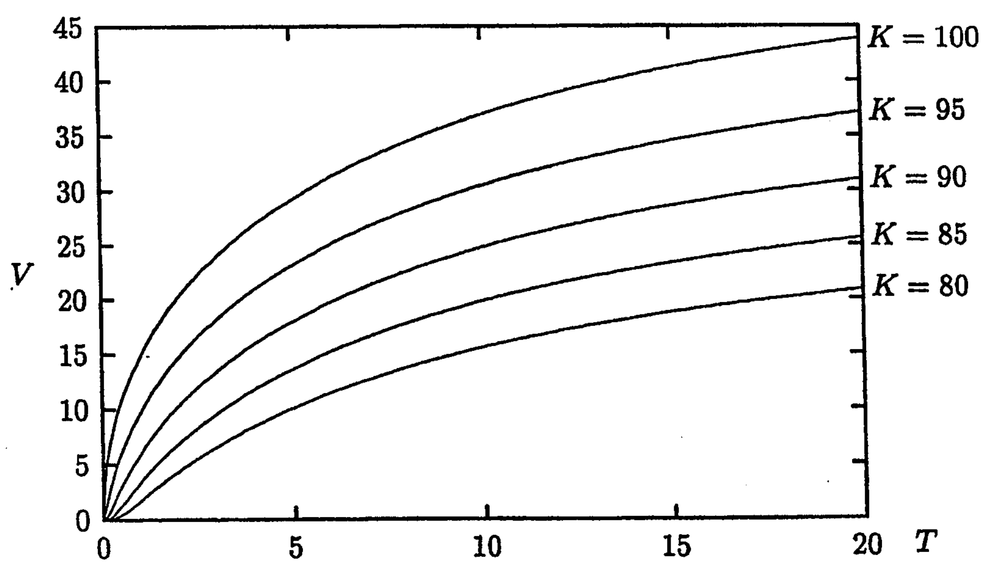

Figure 3

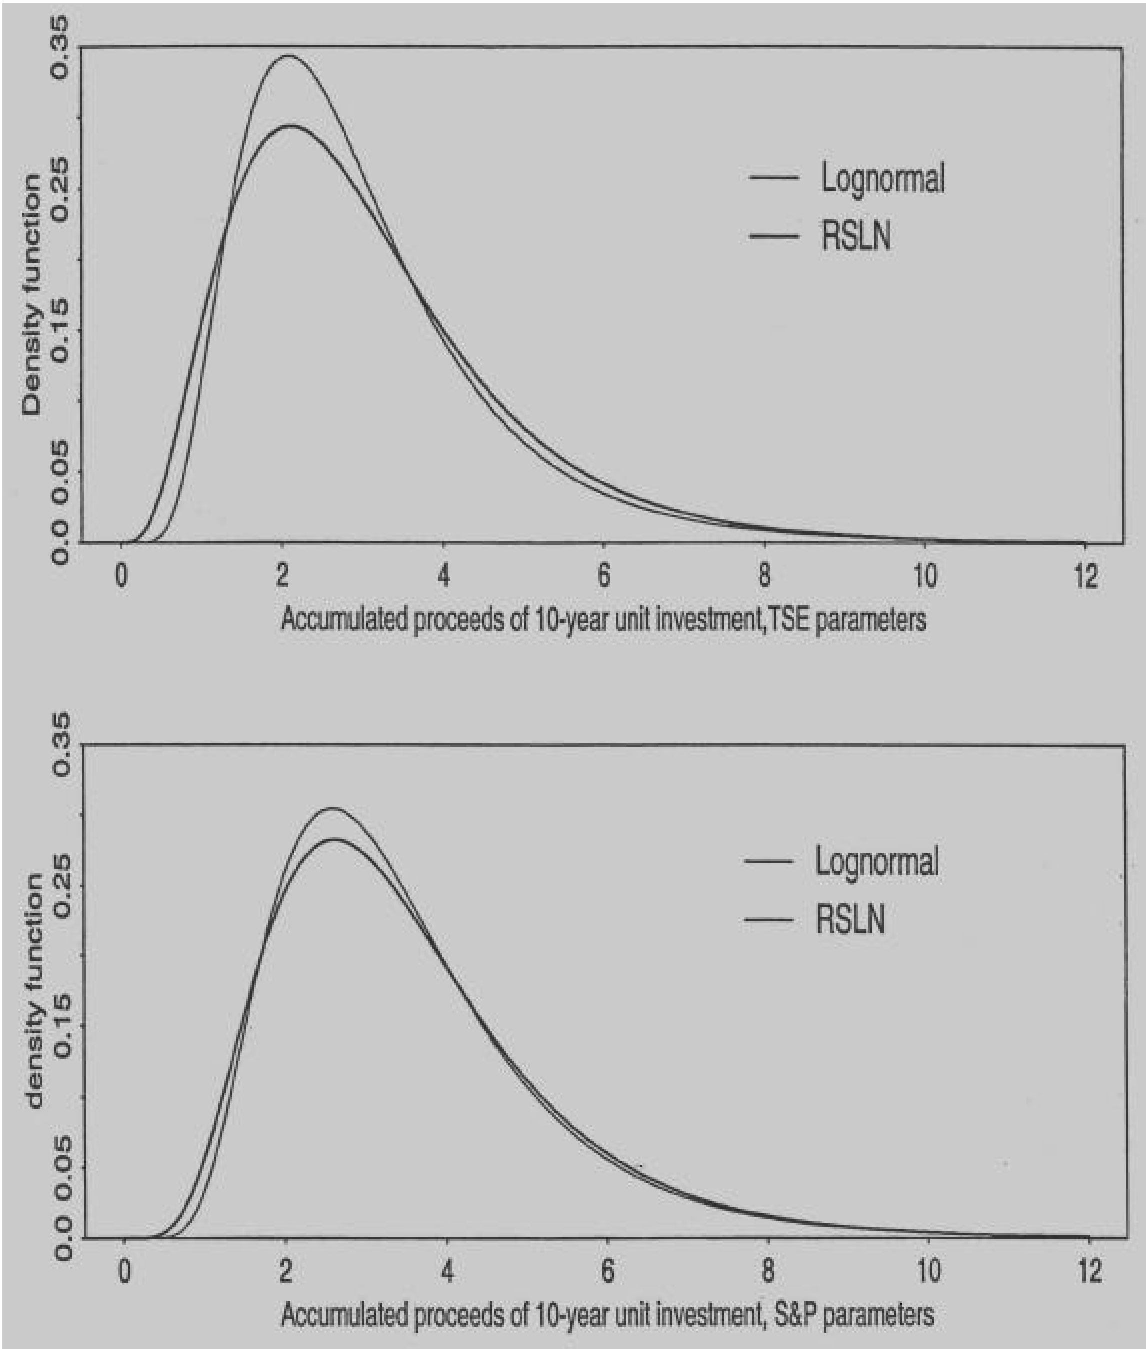
Density function
0.0 0.05 0.15 0.25 0.35
Lognormal
RSLN
0 2 4 6 8 10 12
Accumulated proceeds of 10-year unit investment,TSE parameters
density function
0.0 0.05 0.15 0.25 0.35
Lognormal
RSLN
0 2 4 6 8 10 12
Accumulated proceeds of 10-year unit investment, S&P parameters