RECORD, Volume 24, No. 2*

Maui II Spring Meeting
June 22–24, 1998

Session 79PD
Introduction To Value-At-Risk

Track: Investment
Key words: Investments, Risk Management

Moderator: EDWARD K. MOHORIC
Panelists: DEBORAH K. ORLANDO†
MARK ABBOTT‡
Recorder: EDWARD K. MOHORIC

Summary: Value-at-risk has received a great deal of attention in recent years as an effective risk management concept. This session examines the basics, including:

- *What is it and how does it work?*
- *How much and how is it being used?*
- *What are the pros and cons relative to other risk management techniques?*
- *What are the likely implications for pension plan funding and investments?*

Mr. Edward P. Mohoric: Deborah Orlando is going to be our first speaker. Deborah is a partner with Ernst & Young's Risk Management Group. She'll be talking about the basics of value-at-risk (VAR).

Our second speaker will be Mark Abbott from BARRA, who directs their consulting and marketing efforts, including risk management, asset/liability modeling, custom research, custom valuation modeling, and convertible bond valuation. Mark joined BARRA in 1997 through their acquisition of Global Advanced Technologies (GAT). I'm sure you're all aware of GAT, which provided software, research and consulting to the fixed-income market. Mark's background includes statistical analysis for Merrill Lynch, adaptive stochastic processes research, and applied cancer research. Mark will be giving some VAR illustrations and more details after Deborah's speech.

Ms. Deborah K. Orlando: I'd like to go over the basics of VAR, and then talk to you about how you can apply it in your daily business of managing your investment

†Ms. Orlando, not a member of the sponsoring organizations, is a Partner with Ernst & Young's Risk Management Group in New York, NY.

‡Mr. Abbott, not a member of the sponsoring organizations, is a Director of Research at BARRA, Inc. in New York, NY.

portfolios. How many people do manage investment portfolios on a daily basis? How many people are on the insurance side? And how many people are on the pension side? There seems to be a good mix in the audience.

Before I go into the presentation, let me tell you a little bit about what we're seeing among our clients on the insurance and the asset management side. We deal with a lot of life insurers and property and casualty (P&C) insurers. The life insurers tend to be looking at VAR with some interest. Many life insurance companies have recently obtained executives from the banking industry, and these banking executives want to be able to get the same kind of risk reports they used to get all rolled up into a nice neat package for decision making at the banking institutions.

What we find with many of our life insurance clients is that they're pretty decentralized. Each business unit, and the investment department calculates its own risk, but they really don't roll them up into anything that's uniform and consistent for the management of the company. The executives, the CEO, and the chief financial officer get reports. They get one thing from annuities, one thing from pensions, and so on. They really don't have anything that puts everything in one framework.

We've been talking with them and working with them to build a framework in which they can use VAR concepts in more of an insurance context. We've been looking at things like earnings-at-risk, surplus-at-risk, things that are important to them that allow them to roll their liabilities into the picture, and not just look at the asset side. That's what's going on in life insurance. They're at the very beginning stages of this process. There are some companies that have done that, and a lot of companies that haven't even touched on it yet.

On the P&C side, we're seeing a lot of interest in dynamic financial analysis —simulation techniques to simulate different economic events on different portions of your business. We're working with clients to try to help them take the results of dynamic financial analysis and roll them up into a related risk reporting system using earnings-at-risk or surplus-at-risk. In the insurers that I deal with, there are some who are starting it. There are some who are ahead of the curve, using mostly variance/covariance. Very few are using Monte Carlo yet.

On the asset management side, they're pretty much dealing with assets, so it's not a hard concept to apply. They take the concept straight from the banking world, basically VAR. Many are pretty sophisticated and use Monte Carlo, but some people don't use Monte Carlo. They use the variance/covariance approach instead. It really is a matter of computational efficiency and how much accuracy you need in your reporting to your management. But for some managements who haven't been getting rolled-up results at all, variance/covariance works quite neatly. You make

some simplifying assumptions and it works quite well. For other managements who are more demanding and want to know every single case that could possibly happen, you have to go to Monte Carlo even though it's going to cost you a lot more. That's just some background information.

What is VAR? VAR really is a methodology, or a series of methodologies, that allow us to express risk in probability-based terms. As we discussed, the banking institutions were the primary originators of the VAR concept. Their trading managements wanted to know how much risk they had at a particular point in time. Back in the good old days, in the early 1980s, we used to make up limits for trading and for counterparties based on notional amounts, and that was some estimated percentage. It could be 20%, it could be 10%, it could be 50% depending on the complexity of the transaction. You basically were assuming that your losses would be no greater than that over a certain time horizon.

But then, things got more complicated. Instruments got more complicated. Different types of exotic options came into play. Volatility in the markets started increasing, and people wanted to know more accurately what kind of risk they were taking. Plus, traders wanted to take more positions. In order to take more positions, they wanted to prove that they had a lot less risk than they were estimating they had by just the simple numbers that they were using. They wanted to be able to expand their usage of their limit, so they said, "We'll get more scientific and we'll do more probabilistic-based calculations, and we'll figure out how much risk we really do have." So they started looking at what the standard deviations of the returns of the various assets in my portfolio, or the various instruments in their portfolio were, and how correlations among these instruments interacted to either increase the risk in their portfolio or reduce it through diversification. That's where these concepts developed basically.

Many people used key rate durations (KRDs) and other kinds of duration concepts for their investment portfolios. If you're dealing with equity portfolios, you look at your betas and other equivalents. If you're doing foreign currencies, you may look at spot equivalence of various amounts of currencies. But if you try to roll all that up, it becomes very difficult because they're all on different bases, and that's what VAR tries to accomplish.

Some of the kinds of statements that you can say to your management if you use VAR are, "There's a 95% chance that my portfolio will lose no more than $65,000 in the next 30 days." Or, "There is a Y% probability of earning an X% return over the next year. So it's very probabilistic. There are some caveats though when you go to your management with these kind of numbers. You have to express to them that you're looking at things under a certain specified type of distribution, and there are very fat tails that you might not be getting in some distributions.

We did a lot of work in emerging markets back in my old job when I was a banker. Trying to estimate what was going to happen to an emerging market's portfolio, i.e., bonds, was a very difficult task. You were trying to figure out whether things were going to behave the way they did in a crisis situation, such as the Mexican peso crisis, which would be the real tails of the distribution, or whether they're going to behave as they were beginning to do in the early to mid 1990s when things were beginning to get a little better and a little smoother.

In my last few years as a banker in 1995 and 1996, you basically saw correlations coming together—less diversification effects, correlations coming closer to one, and volatilities going down. So if you looked at recent events, you would say, "Well, you know, things are getting smoother in the markets. They're not that bad in emerging markets." We were trying to determine margin requirements so that our traders could deal with their clients. If you go look at the Mexican peso crisis and you base your margin requirements on what happened in the Mexican peso crisis, your traders are never going to do any business because the margin requirements are going to be so high. So you have to balance that out. Anytime you give a VAR number, as nice as it sounds, you have to explain to your management what that really means. Given this assumption of this kind of a distribution, there's a 95% probability that we're not going to exceed this kind of loss, but there's a 5% probability that we could. It's also good to do stress testing along with a VAR number because you want to look at the extreme causes.

From the Floor: When you say those words about the distribution, do the eyes of the people you are talking to glaze over, or do they absorb what you are saying?

Ms. Orlando: I worked with a lot of different boards of people. Most of the people that sit on boards of banks or boards of insurance companies are pretty smart people, but they come from different backgrounds. Most people understand the concept of distributions, but they understand more a blanket statement that they won't lose more than this in a certain amount of days. But you have to tell them, "Well, just be careful because there's a 5% chance that you will." I would explain to them there's a 5% chance you could lose more than that. I would explain it to them in terms of distributions. I'd say, "We're basing this on the way the events have been happening recently; not on a crisis event like the Mexican peso crisis." They can understand that.

Let's discuss how VAR was developed. VAR was developed for two purposes: estimating losses and allowing the calculation of risk-adjusted measures. Obviously, if you're a portfolio manager or a trader, you can create a lot of return by taking a lot of risk. Financial Institutions are trying to actually apply it in a lot of insurance companies in terms of risk-based capital (RBC) (using VAR concepts for more accurate representations of RBC). They are trying to figure out what business

lines are producing the best return for the least amount of risk. It's similar to the idea of a Sharpe ratio, where you're looking at your return for a unit of volatility.

You're looking at how much return you're producing with what amount of risk. It gives a better representation than just looking at the amount of return that a particular portfolio or particular business line is producing, so it could be used for capital allocation or asset allocation.

There are two different big-picture approaches to VAR. One is parametric VAR, which basically looks at the standard deviation of returns of the portfolio, and the full valuation VAR, which generates distribution. The Monte Carlo that we talked about is an example of a full valuation methodology, and J.P. Morgan's Risk Metrics variance/covariance type approach is what they call a parametric value. Each one has its advantages and its disadvantages, and we'll go into that.

First, let's discuss parametric variance/covariance. You first take your positions and assign them into risk buckets. You do cash-flow mapping for your fixed income. You may do delta weighting for your derivatives and beta weighting for your equities. Then you either prepare a covariance matrix or you download one. You can go to an outside supplier and get a covariance matrix. What the covariance matrix does is show the relationships among the different types of asset classes.

If you're looking at currencies and equities, this would be how the Standard & Poor's (S&P) 500 is or isn't correlated with the movement of the dollar yen. It would also be how Treasury bonds are or aren't correlated with the movement of the S&P 500. All those factors get put into a big covariance matrix, and some of these covariance matrices that you get from a data supplier are very large. They show all the different relationships. That also goes into your calculation of variance/covariance. You calculate your portfolio volatility, and then based on the confidence level that you want to choose, you would basically apply a multiple to the absolute value of the standard deviation of the portfolio.

For example, if you wanted to calculate two standard deviations with a 97.5% confidence level, you would multiply the portfolio standard deviation that you got by two to come up with what your VAR would be times the notional amount of your portfolio. So variance/covariance is a pretty easy methodology to apply except for the covariance matrix, which you either have to develop yourself or get from an outside supplier.

If you use variance/covariance methodology, you're assuming a normal distribution. So you're not assuming anything with fat tails or anything that's skewed. You have to do that in order to assume your confidence levels and the number of standard deviations that you're going to need. It is the most commonly used methodology

today. As I said, it's pretty easy to apply, and pretty reasonable in terms of the cost. It has several drawbacks because you are making some simplifying assumptions, and we'll talk about those drawbacks in a little bit.

Let's look at what has been happening to correlations in the recent past. Chart 1 shows the correlations between U.S. stocks and Treasury bonds: These are weekly observations over rolling one-year periods going from about December 1990 to January 1997. As you can see, the correlations have been quite volatile. They can be as low as -1 or as high as +0.7, and sometimes they're zero, so it's very, very difficult to take a static view of where the correlations among different instruments are because, as time goes on, they change, and they've been changing more frequently.

Now let's look at what has been going on with volatility. Anybody who has invested in the stock market with your personal portfolio can appreciate what has been going on with volatility of stock prices recently. Chart 2 is from an earlier period, again, April 1990 to June 1997. As you can, see stock prices have been quite volatile. Since volatility is moving a lot (even more so recently) and correlations are moving a lot, it's very hard to make simplifying assumptions that are static. Let's look at the full valuation methodologies and then we'll compare them. Actually, let's sum up the good things and the bad things about variance/covariance. The good thing is it's very, very easy to apply, relatively speaking, if anything could be easy to apply here. You don't have to build tremendously big systems. You can do it on spreadsheets. We did a lot of work when we were first going to implement value added risk (VAR), at one of my old firms: We did it all on spreadsheets. You don't need a tremendous amount of computational power.

The bad thing is that you're assuming normal distribution. If the assets that you're dealing with are not normally distributed, it could be a problem. You're also not taking into consideration second order effects like gamma. You're not looking at optionality, so it's very difficult to do VAR calculations on options using a parametric approach or a variance/covariance approach. Those are the two main drawbacks.

If you look at some full valuation approaches, there are two big categories; historical simulation and Monte Carlo simulation. Historical simulation basically looks at what happened in the past? I look at the historical market movements that happened over the time period that I have the data series for, and I assume that those historical market movements will repeat themselves at the present time with a couple of caveats. I look at my current portfolio allocation rather than what it was in the past, and I also design specified pricing functions to take care of nonlinearities. Now I can price options. I can do VAR on options, so I'm actually specifying pricing functions.

Then I generate a theoretical profit and loss based on those historical market movements. I calculate distribution of returns, and then the observations are ranked, and the desired probability level is chosen. Again, I decide if I want 95%, 97.5%, or 99% confidence. I can show all these results via a histogram.

Historical analysis avoids some of the drawbacks of the parametric approach of variance/covariance because it allows me to take into consideration non-normal distributions. It looks at all the observations. It also allows me to take care of nonlinearity. I can actually do VAR on options using historical simulation. However, there are two things that are not so good about it. I assume that the present and the future are going to behave the way the past did by looking at the way things moved in history. Second, I must have a large database because where am I going to keep all these numbers that I must use to calculate and to look at what happened in history? For some instruments, particularly emerging market instruments, you don't have a very long-term database. You have a three or four-year database. You have to look at the particular instrument, and you have to decide, for that instrument, whether it's better to use a variance/covariance approach, or whether it's better to use a historical VAR approach based on what you have at your disposal.

What we're doing in Monte Carlo simulation is generating random paths for each of the market factors. We're looking at all the 10,000 iterations, and what could happen to that market factor in the future. Basically, we determine gains and losses for each instrument. We sum up the results, and we produce a probability distribution. We rank the results and show them via a histogram. Or, we determine what confidence level we want so we can determine what our VAR is.

From the Floor: How do you work in so many moving parts in a simulation?

Ms. Orlando: You hold everything else constant. You just vary one factor. You can't move everything all at once.

From the Floor: And you relate everything to that one factor?

Ms. Orlando: Exactly. Then you turn around and do it holding everything else constant and just varying another factor. You then create your variance/covariance matrix. As long as you're not double counting things, as you said, you really do come up with a pretty good result that takes into consideration most possible outcomes.

Mr. Mark Abbott: In terms of the horizon that one is using, VAR assumes you're doing nothing. So you're going to simulate no response to deteriorating capital, and that's not really realistic either. With too long of a horizon for VAR, unless you're

really limited where you can't do anything—something that may not be considered. So we're looking at reasonable horizons, in terms of some plan of action, and the assets and liabilities. Banks typically look at these very short because they're trading positions. They're going to be able to liquidate positions and book losses right away.

From the Floor: You used some sort of probabilistic model or something for Monte Carlo simulation. Isn't that based on historical information?

Ms. Orlando: No. It's based on what you think is going to happen. Then it generates random paths, so it takes into consideration everything.

From the Floor: But wouldn't you use historical information to decide what you think is going to happen?

Ms. Orlando: You might, or you might just be a genius.

From the Floor: In terms of using a model for the future, it could be a mixture of a model that is based on past data. Also if you take a financial economic approach, then you're also taking into account just what information you have in today's market.

Ms. Orlando: It's a combination. If you're using a parametric approach, you're assuming your portfolio stays constant. You're assuming you do absolutely nothing to fix it, and you're also assuming that your deals don't run off. For example, if you're holding a five-year Treasury bond and time passes, now you're down to four years. It assumes it doesn't become a four-year Treasury bond, but it still stays a five-year Treasury bond. That's one of the simplified assumptions with variance/covariance. Whereas in Monte Carlo, you can assume the aging of your assets. It actually takes into consideration that your assets are changing and becoming shorter as they go along.

Monte Carlo simulation has some benefits because it doesn't assume a normal distribution. It also can be used for options. It allows you to age your portfolio. It's considered the most accurate because it takes into consideration the most possibilities, and it allows you to look at the details of the distribution by doing a random generation. It's very complicated in terms of computation, and you need a lot of computer power for it. With very few exceptions, you can't really build Excel spreadsheets big enough to hold your Monte Carlo simulations.

Chart 3 shows sample output for VAR calculations using Monte Carlo and variance/covariance methodology. By using the two approaches, we've actually plotted the histogram of the occurrences and we overlaid a normal distribution

curve, which is what you would be looking at if you're looking at the variance/covariance methodology. Here you can see the two VAR calculations.

The gray area is the 95% Monte Carlo VAR, which comes in at $592 of VAR. For the variance/covariance, you basically go to 1.65 standard deviation, exactly where it would be on the normal curve. That comes in at $550 of VAR. You can actually see that you do get different numbers. In this particular case, it's only a $42 difference in VAR.

Now with this representing $42 million, that may mean something to you. If it's $42,000, maybe you could care less. Maybe it's not worth investing in the system that you would need in order to produce the Monte Carlo VAR, but that's the way it would be derived. The variance/covariance is simply derived from the normal curve where your 1.65 standard deviations would be for 95% confidence. It's a pretty easy concept.

I want to talk briefly with you about some applications. If you're going to develop a VAR system, or an earnings-at-risk system, or a surplus-at-risk system, what do you need to have in the system? You need to have your positions, your assets and liabilities, your quantities, and descriptions of them. You need to know your market data. You need your curves, your currency rates, your volatilities, your correlations, and you have to map your market factors. You have to maintain your static data, and you have to generate your distribution of returns, which we've done in the early discussion. You need customized risk reporting. What can you get out of this VAR system? You can get VAR or earnings-at-risk or surplus-at-risk. You can have a risk disaggregation which shows what parts of your business contribute what risks to your business. You can actually disaggregate your risk and create "what if" analysis.

But if you're an investment manager and you have a portfolio and a specific asset allocation, you can calculate your VAR of that portfolio as it now stands. Then you can ask, "What will happen to my VAR if I take out a particular manager, or if I add a particular manager? Or, if I take half of my assets that I have in equities and convert them into bonds, what does that do to my VAR?" You can do all those calculations, so when you're making your asset allocation decisions, VAR can be a pretty helpful tool.

The same thing if you're an insurance company. You can look at not only the risk by business lines, but also the risks by type of risk: interest rate risk, mortality, morbidity, lapses, credit risk, currency risk, equity risk, event risk. You can look at the amount of risk that each type of risk contributes to the total risk of the portfolio, and the amount of risk that each business line contributes. Let's say your annuities line wants to launch a new product, and that product is going to add a substantial amount of credit risk to the total risk of the company. If you're looking at that line

just in isolation, you may determine that you shouldn't launch that product. But if you look around the whole company, you might see how the rolled-up risk looks when you bring up all the risks from all the different product lines. Maybe there's a product line where there's a product that actually diversifies that credit risk that you would be taking on in the new product. That's why it's so important to be able to roll up your risks because you're not as efficient as you could be if you don't look at all the risks in your whole company. You can't just launch a new product and not worry about what it does to the total risk of the company; however, you don't want to prevent yourself from launching a new product if there are diversification effects. That's what VAR allows you to take into consideration.

If you are going to do VAR analysis, we always recommend that you do some stress testing on your portfolio. It's very, very important because when you're looking at probability distributions, or you're looking at VAR and you're saying, "Oh, I'm 95% confident, or I'm 99% confident, or I'm whatever." Based on the various assumptions you're making, you're missing something. You're missing either the tail, the 5%, or the 1%. You're basing things on historical data, you're assuming a normal distribution, or you're doing something.

You always want to do some scenario analysis to look at the absolute worst cases. Those should come about by sitting with your management and saying, "What's the worst thing that can happen to us? Interest rates drop below 4% and all our minimum guarantees come in. We can't invest profitably." Now I can determine what can happen in the worst case scenario. You have a number now that you can use to go to your management group and say, "I'm 95% confident. However, if everything goes wrong and interest rates go down below 4%, this is how much we could lose." So you could also add that, and we always recommend it.

Here are some applications of VAR. These can be for investment managers. These can be for anybody—investment managers in an insurance company or pension fund or whomever is trying to make asset allocation or capital allocation decisions. What can you use VAR for? You can use it to calculate how much your risk is reduced or increased by applying different strategies or investing in different asset classes. Or you can measure risks associated with different regions, different managers, or different product lines and determine the probability of obtaining a specified level of returns. You can use it for benchmarking. You can use it to determine if your benchmarks are correct on your various portfolios, or if they're reachable, because maybe you have a benchmark and maybe it's not reachable at all, given the way the markets have been moving and the way volatilities have been in the market. Maybe you'd have to take too much risk in order to get the kind of returns you need to meet your benchmarks.

You can determine potential losses associated with individual pension fund managers, and you can also use it to evaluate risks dealing with various counterparties. That has been a big trend in the industry. More and more people are focused on credit. More and more people are focused on how to better calculate their credit risk rather than just pull some number out of the like 20% or 10%. People also are starting to use up their lines that they have with various counterparties, so they want to be able to do more business with them. If they can get a better hand on credit risk by using probabilities and volatility and figuring in correlations, then they can do more business. That's pretty important. In terms of slicing and dicing risk, you can decompose your risk into different types, like credit risk, interest rate risk, equity risk, different business lines, asset allocation, or hedging decisions.

By using VAR technologies, you can look at how effective your hedges are going to be on your portfolio. You can be more accurate. Instead of just using sensitivity-based measures like delta, gamma, vega, theta and rho, you can actually calculate VAR of both the hedges and underlying instruments in your portfolio and make sure they match. If you add a hedge to your portfolio, how much does that reduce your VAR.

As I told you, more and more people are getting interested in risk-adjusted return analysis because they want to allocate their capital more efficiently, so they want to make sure that their managers or their business lines are not just creating return, but doing it in the most efficient way possible with respect to risk. Short-fall probability. What's the probability that I'm going to miss my target or miss my benchmarks? Benchmark divergence is pretty similar. What's the risk of my portfolio versus my benchmark? That has to do with tracking errors. There are a lot of different things that VAR can do for you.

What are the system alternatives that are involved with VAR? You can develop your own system. If you're going to develop a variance/covariance system, you can probably do that internally. Then get a data outsource to give you the variance/covariance matrices, or use something like Risk Metrics. You can purchase a system if you want to do something a little more detailed and closer to the full valuation methodology. There are several commercially available tools. Most of the systems start around $300,000 for those we would use in a bank trading environment. Or you can outsource it through financial institutions. If you have some good counterparty relationships, you can ask them to calculate your VAR for you. You can also use consulting firms for that. So those are your alternatives.

Mr. Larry H. Rubin: One issue we've been dealing with is we find that our measure is very sensitive to the stock risk premium assumed. How do you deal with that in a presentation to management?

Ms. Orlando: You have to really do a stress test on that. You calculate your numbers and if something happens with the stock risk premium, then you can tell what is going to happen. Any factors that are really, really sensitive, or that don't behave in a relatively normal fashion, will have to be given a stress test.

Mr. Andrew Kearns: In terms of the issue of deciding between variance/covariance or Monte Carlo, the reason why I personally use Monte Carlo is because I'm looking at pension funding problems where you're looking at a longer term time horizon. In the short term, there is a reasonable degree of symmetry, so the variance/ covariance works quite effectively, whereas, as you mentioned, in the longer term, nonlinearities come in. This is an instance in which I feel Monte Carlo is almost necessary to make sure that you get the right shape of the distribution because if you're looking at the tails of the distribution, and you assume symmetry in normal distributions, then you're going to be way out.

Ms. Orlando: Exactly. That would be the same for insurance companies. They're looking at very long-term liabilities. It's the same kind of situation.

Mr. Abbott: I'm going to talk a little bit about some of the same issues, but from a slightly different perspective. I think we both have come from more of the investment side, so we focused on a lot of the risks that investment managers have focused on—the interest rate risk.

BARRA certainly has a lot of factors in this variance/covariance matrix, including global equity risk factors as well.

Risk for the different credit qualities and prepayment risk for the underlying collateral of mortgage-backed securities is needed to enhance the matrix provided by J.P. Morgan in their Risk Metrics data set. So it isn't sufficient but you can really get a good handle if you're using a parametric methodology. What I want to do essentially is to cover four areas. In valuation, using a parametric approach, you need to have arbitrage-free framework so you can do economic valuation. That is going to work for both assets and liabilities if you have liability models, and you're dealing with economic valuation of those liabilities.

Specifically, I wanted to look at some examples. We build up our delta-normal parametric approach for VAR using the delta methodologies of shocking various risk parameters. We go a little bit further on the interest rate side than the J.P. Morgan methodology. We use key rate duration to look at the non-parallel interest rate exposure, which explains some of the embedded optionality. I'll explain that later.

Explaining the optionality of the convexity or the gamma can enhance the delta-normal method. I prefer this delta-gamma method for parametric VAR. Again, that

works within a relatively short horizon. I'd extend that for an insurance company or pension plan. They could perhaps look at it over a year. Maybe five years would be the maximum, but certainly, you should look at it longer. So the simulation approach is appropriate.

Typically, we're using valuation models where we are doing valuation over many different paths. For those securities or liabilities that have long exposures where you really can't do anything for that period of time, you probably already have a fairly rich start in the universe. I'd say you probably need at least several hundred or maybe thousands of scenarios for those when you do a Monte Carlo simulation. The problem with Monte Carlo is as you get more and more risk drivers, you really need to have more and more paths, and it becomes unattainable if you're dealing with a complex multi-currency, multi-factor liability and multi-factor asset model universe. It can give you some sparse understanding, and you can fit the distributions, but it doesn't give you the confidence level, I think, that you would require even for a large number of scenarios.

I'm going to take the VAR methodology from the parametric approach, then apply it to VAR of surplus, and then build that up to a corporate model. Deborah touched on some of those aspects, in particular earnings-at-risk. Essentially, you need to look at all of the factors from the shareholder's perspective if you're dealing with the corporation's value, or even maybe the pensioner's value, depending on the perspective or the client you're actually addressing. I'll talk about some of those applications.

If you're dealing with a fixed-income world, you need to have an underlying model to estimate fair value. The factors of the valuation model and the need for calibration must be addressed before building a covariance/variance matrix. You have certain assumptions that determine the starting point of your simulations. One way is to calibrate to the market. Certainly, you could also calibrate things to past historical periods, i.e., look at starting points that were in the past, and generate covariance matrices based on actual periods of data rather than just the most recent data.

In estimating the market value, which is the starting point in parametric methodologies, we want to get a price, shock some risk parameter, and see what effect that has on the price. We need to build up a framework of a spot curve just to get at the discounting effect of the discount rate, an arbitrage-free lattice that calculates embedded optionalities, and path-dependent models to determine the price sensitivity of structured products to the prepayment effects of mortgage-backed securities.

You might also need to enhance this model with movements of inflation equity and prices. All those components may need to be built together in the case of very complex structured notes, convertible bonds, or even liability pension plans where you're looking at the sensitivities to all those different factors.

I first wanted to provide a general definition of duration that is not confined to interest rate risk measures. Duration is essentially the price sensitivity when you are shocking one parameter. In this case, it would be the effective duration interest rate for fixed income, when you're shocking rates in a parallel fashion. So as rates rise and you have a fixed cash flow, its present value is going to decrease. That would be a positive duration. Rates rise when the price goes down. For equity options, this sensitivity to the underlying stock price is called delta.

You get a negative duration if rates rose and the value increased, as you might have in a leveraged index security. That would be termed negative duration. With options like calls and puts, you can see very large positive and negative durations. Same thing with caps and floors. Leverage essentially amplifies the positive or negative duration.

If the yield curve doesn't move in a parallel fashion though, you need to worry about a sensitivity to a non-parallel shift. If the yield curve twists or kinks, or if you're just looking at specific risk to a particular key rate term, you should be using a methodology called KRDs. The beauty of this is it actually can be used as a hedging tool, especially if you're trying to replicate or target an index. If you're trying to hedge your portfolio against these movements relative to an index, we found that KRDs give you a much more stable solution than just using duration. That is because duration is not really constraining your portfolio against movements that are not parallel.

Beyond that, we apply interest rate movement simulations for the option adjusted spread calculations for mortgage-backed securities. We apply simulations to value liabilities for the corporation, bank, life insurance, and the annuity worlds, and, of course, for the pension plans that we're working with. We're dealing with multi-scenarios and some fixed costs and an appropriate spread that's associated with that.

But there is tremendous information in all of those simulations, so I do agree that it's good to look at both sides of the equation—parametric and simulation—because together they tell the whole story. Parametric gets you there faster with less cost. Then you can maybe spend your resources looking at those areas where there seems to be a lot of risk to better explain it.

This is an abbreviated description of the method to compute KRDs. Basically we start with an initial yield curve calibrated from market prices and a small normal

parallel shift for duration. Next decompose this full shift into partial shifts. Chart 4 has an example of one in the middle around the 15-year point from 10 years up, and then down from 15 to 20 years. We would calculate three prices: a price along the unshifted curve with the partial shift up, and with the partial shift down. From those prices we can calculate duration and convexity to this key rate term shift, if you will.

Then we can sum up all of these partial shifts and we get a parallel shift. In fact, you find that the sum of the KRDs equals the effective duration. This is true except where there are binary options where, for example, something has been triggered beyond a level and you're not getting that level in the partial shift diagonals. But by averaging over different shifts, you can include that value as well. So there are techniques you can use to make this methodology robust for all cases.

For example, a 20-year zero-coupon bond has no KRD sensitivity to any of the shorter rate movements because essentially the discount rate effect is unwound before it hits this cash flow, if you will, at the 20-year point in time. But all the effective duration is essentially realized at this one point.

If you're dealing with a security that has embedded options, as you would here in this corporate bond example in Chart 5, the gray represents a premium coupon to the current rate. You can see that it has a lower sensitivity where there are less cash flows that would be sensitive to interest rates.

The optionality pushes it a little bit shorter, if you will, in terms of where the effective duration, or the KRDs are being kicked in. If you're actually looking at annuities, or caps and floors, or calls or puts, then you get a mixture of positive KRDs, and offsetting negative KRDs. They're quite leveraged. Twists of the yield curve can really change those values tremendously, so these are the things that you need to be aware of as you're starting to develop a parametric model.

We actually used the KRD terms associated with the volatilities and the correlations on the underlying spot rate movements to build up our covariance matrix. We find that actually gives fairly stable results even for the delta-normal methodology. That is what you would use as the classic equity risk terms and the corresponding risk measures that you might have for a fixed income. How we can actually start to put these together in a VAR sense.

In terms of market risks, if you're dealing with VAR, it is fairly difficult to deal with all of the asset classes to build this covariance matrix. You're dealing with complex liabilities where you have human behavior, people moving from one type of fund to another. If people are actually leaving the company and changing the funding

aspects, you have to model things like behavioral risks. Obviously, as you're building up your models, you have model risk.

If you're dealing with funding issues, then you certainly have even more complex long-term issues and asymmetries that you need to consider. Maybe adding a convexity term and using a parametric delta-gamma methodology is more appropriate. That becomes a lot more costly in terms of computer resources and the calibration.

Deborah provided a good definition for VAR. I want to elaborate on several things you need to consider. Your starting set of assumptions is very important. VAR is the maximum expected loss for some set holding period, and a certain level of confidence. Those are the driving forces, and, obviously, all the other risk factors, the model risk and the data that you're using to calibrate this, are critically important. Regarding the assumptions of your distributions of volatility, even if you're dealing with models where you're looking at projecting, you might calibrate this to actually a little bit biased in terms of your expectations for returns rather than the set market values. That's something that you should investigate and can change. You can recalibrate the current market conditions based on those hypothetical horizons, and look at VAR at those horizons in the same sort of way.

We're finding that VAR is typically being used in monitoring and measuring risk in the trading and portfolio management areas. Only one client that we worked with so far has actually reported this in their annual report, but I think it's going to be something that will happen more often in the future. BARRA actually acquired Rogers Casey two years ago, which is a plan sponsor consulting firm. They just released a risk report where they surveyed the importance of risk and all of its various attributes across a wide sampling of corporations. That's a very interesting study.

There's a total plan risk in which you're looking at active risk to the indexes against which plans are measuring their asset managers. Looking at the risk exposure is something we've also been working with, both at the sponsor level and at the custodian level in terms of reporting your VAR. I think you're all going to see a lot more of VAR in the near term. Probably the parametric, variance/ covariance, matrix information will be dominant. In that information, the methodology is to extend this duration and shocking concept to all the different risk drivers. We're going to look at this in terms of the information that comes even at the per instrument level, or you can look at the liability in terms of different segments of the liabilities. You can drill down and look at the risk exposure that you have either on the asset side, the liability side, or at the surplus level, or you can roll it all the way up to the top tier corporate level or segments of the market, different asset classes, or different portfolios if you're looking at funds of various sorts. If you're looking at

the funds of funds equation, you can really see how well the diversification of the funds of funds is adding value.

In the model, you have your different risks, and you have your different holdings, like asset and liability holdings, and a distribution. We're calculating the loss at a 1% level. Certainly, the distributions may be skewed, non-normal, or lognormal. You need to think about how you're going to handle those down the road if you use the parametric approach.

We want to calculate the VAR, so we must look at the total VAR across all of these different terms. We actually look at the different KRD terms and the correlations of those KRD terms versus the different asset holdings. Then we build up this equation to calculate VAR.

The correlation matrix. In an example across all of the different key rate exposures, I've added the mortgage-backed security, and some of the credit quality sectors. It gets quite large. If you're dealing with multinational diversification, this can get very unwieldly, and you may have to simplify your exposure to more aggregated terms to get some results.

From the Floor: When you're looking at the triple A, are you looking at correlations for the spread over Treasuries?

Mr. Abbott: Yes. We're looking at the additional spread correlation point when performing the variance/covariance matrix calculation. Fortunately, we have a tremendous database of the historical spreads. We even have option adjusted spread for different tranche types of collateralized mortgage obligations (CMOs). We're doing a lot of VAR calculations for clients that have large exposure to CMOs. We drill down to a very detailed level with the data that we have available to us. VAR is usually considered at the aggregate level, and you can, certainly, calculate it at the security level. You want to look at risk relative to the contribution to the total considering the diversification effect from the covariance matrix.

Once you've calculated the VAR for the whole portfolio, you also want to look at how the individual VAR of any asset correlates with either your sector VAR or your total VAR. Then we can look at the dollar contribution to total VAR. We call that dollar beta, and that's an additive measure.

Risks can be reported relative to the total VAR. You can look at the contribution of risk by each item. Or you can look at the dollar beta over the total market value to get at what BARRA has always called the marginal risk contribution. Before VAR was actually the buzzword, BARRA was reporting this information. It started actually when VAR was more known for equity risk exposure in factor models than

fixed income. But since BARRA acquired Global Advanced Technology, it is a much more dominant force. On a worldwide basis, we're making quite a push to get all the assets models to integrate globally for one consistent way of viewing risk. You want to sort of slice and dice, as Deborah mentioned, and get at the very sources of risk. At the individual security level, you might have to identify securities that have extremely high contributions, so you look at the dollar beta relative to a particular sector or the total portfolio. The nice thing about some sort of a table is you can look at the different asset classes and your liabilities, whatever they happen to be.

You need to have the risk sources that are appropriate. In the different practices, health is going to have several uncorrelated or very correlated risk sources that aren't in this picture. But we would have key rate risk, basis risk, or prepayment risk. We might also have model risk in where you have model parameters that you've calibrated that you want to shock. You can actually then aggregate the VAR by the different business units, and then across risk sources net the exposures.

If you're dealing with surplus, you can look at your total risk exposure at the surplus level using this as well as at the A/L level. By using the diversification dollar beta techniques, we can sort of look at the risk contribution and its relative diversification factor. Where this is going to be strategically important, and I'll get to this a little bit later if I have time in the talk, is how it relates to capital allocation.

If you're dealing with some asset allocation problem, you can actually apply VAR to control and provide information that you would put into your optimization for asset allocation. You should be trying to maximize the risk-adjusted return on your capital, i.e., optimally allocate the risk capital that you're putting into your various investments. You want to maximize the risk-adjusted return, yet stabilize it relative to the surplus, minimizing the volatility of surplus.

Once you see numbers, what make more sense is how the information is going to filter down. Table 1 is an example. I've done it not at the individual security level, but I've brought in, at the aggregate level, the Treasury part of the investments, the corporates, and so on.

TABLE 1
SURPLUS VAR CONTRIBUTION DETAILS

Portfolio	Market Value ($ million)	VAR	$Beta	$Beta/ Total VAR	$Beta/MV
Treasury	$3,628.00	$11.31	$4.52	42.73%	0.14%
Corporate	2,170.00	4.92	-4.28	-40.40	-0.20
Mortgages	625.47	5.47	-4.79	-45.27	-0.77
Loans	1,231.46	30.49	-22.54	-212.83	-1.83
Private placements	2,854.00	33.46	-28.15	-265.81	-0.99
GICs	1,959.94	5.83	3.24	30.59	0.17
Structured Settlements	289.00	1.56	0.98	9.28	0.34
SPDA	443.08	11.69	9.55	90.22	2.16
UL	5,250.00	44.62	36.89	348.38	0.70
WL	1,146.00	19.85	15.16	143.13	1.32
Surplus	$1,078.91	$10.59	$10.59	100%	0.98%

This can also be broken apart into various funds. If you were dealing with a multiple asset manager situation, you could look at each of the asset managers, and you could also then drill down and look at them relative to the different asset, or even the per security level if you needed to. One could look at the contributions of risk to each fund. The same thing applies on the liability side. We typically deal with an economic market value. We're going to look at the liability models and generate various cash flows, and we're even going to do a present value of them. Then we calculate the VAR on each of these, and if you were to sum these up essentially looking at the different sides of the balance sheet, you would not get this number. You would definitely not get the total VAR because this is including the diversification effect through the variance/covariance methodology. This was using just the delta-normal methodology. You're getting these shocks to these parameters, and you're just extrapolating in either direction to see what the risk might be for certain movements.

The model of VAR is taking into effect the volatilities, the current rate levels or parameter levels, and doing the appropriate shocks that are in the appropriate amount to measure sensitivity at the confidence level for the holding period that was specified. Dollar beta basically is going to transfer VAR to an additive measure. If you were to take these, they would sum up to the VAR number. We then can look at the dollar beta, which is relative to the total VAR and to the total market value—the marginal contribution to risk. By looking at these numbers, we see that, in terms of the risk components, we have nicely offsetting risks.

In Table 1, universal life seems to have one of the largest liability dollar beta exposures, and that's being offset by the private placements and perhaps the loans.

In this case, the mortgages seem to be fairly small in terms of the exposure relative to the total portfolio. This was a rather conservative group. But this allows you to exactly pinpoint where your sources of risk are in your aggregated view. You might want to just change this a little bit. If you're uncomfortable with this level of risk, it would be harder to change your underwriting, but certainly you could start to emphasize other product areas more, change the product features to make them more or less attractive, and see how that might even change the benefits. You can look at more complicated funding issues for pensions.

Let's look at diversification. If you were to go into other markets or internationals, how might that further diversify and reduce your exposure? Ultimately, that's something that can enhance value. Our clients use a statistic called the information ratio. The larger you can make the information ratio, the more likely you are to earn greater returns.

I'm going to just jump ahead a little bit and talk about the corporate model. It looks at the straight economic value for the VAR. It becomes more complicated once you start to deal with the tax considerations, the actual capital that you have to set aside in the surplus accounts, and the required surplus that you might need for various purposes. There certainly will be a cost and a change to the VAR calculations that you have. There will be an increase in the level of complexity. We want to take this methodology and translate it to the balance sheet. We'll move more toward the accounting measure so that corporate management can actually get a better handle on risk in terms of shareholder value.

Certainly, we can simplistically think of surplus being the asset value minus the liability value. But what we really want to do is relate VAR to shareholders value. That is the information that shareholders will be particularly interested in. It will tell them a lot about the nature of the risks that are being taken by a corporation, or by an investment strategy. Basically, it starts to address the issues of extremely good performance, but what have the risks been that you've been taking all that time? We need to develop a model that really looks at the various things, such as growth, the tax implications, and the cost of this required surplus—essentially the underlying rate that the company charges for risk capital, and the monitoring of these risk adjusted target returns. Releasing profit may not be an appropriate measure for the pensions. Looking at the shareholders' value may be a little bit different for pension plans since everything is sort of being pulled in together. This is the model that one needs to start with and to adjust appropriately as a model for this release or use of capital. Perhaps you can move it into the investment area after it has already done its job in terms of growing in surplus and freeing up this capital to be invested in other areas.

What this will allow us to do is really manage the various liquidity risks you have if you have particular draws that occur at some point down the line. If you're company happens to be concerned about a downgrading, you can control that likelihood by stabilizing your surplus, and I think that can be a very good thing. Basically, this is the definition of the risk capital. It's going to be the amount of required surplus that we might need, times the cost of capital, divided by the risky rate. That's going to give us the risk capital charge that we're going to have. What we want to do is maximize the return on this risk-adjusted capital, so we look at the returns that we're going to see, divided by the risk capital, and that gives us this return on risk-adjusted capital (RORAC). That's going to be the corporate measure that we're going to need.

Depending on the company, this may change a little bit. You might have built-in expenses to your various models and to the different investments. Certain spreads are allocated to those expenses of the ongoing business. You're going to have certain money that may be profit that's released. Maybe you're going to reinvest that in your pension plan, or maybe set it aside for some future use. Perhaps funding might require future sales. The model is going to look at the risk capital and give you a VAR number on this. I think that's going to tell us the real exposure that the organization has whether they're a life insurance company, a pension plan, a bank, or any other sort of corporation. This model, with a little bit of tweaking, can fit all of those organizations. People at the top have to be responsible to shareholders.

The corporate VAR model is more involved than basically looking at the statements, the present value of the different lines of business and economic VAR analysis. The benefits are that we can also look at the risk allocations that one has. We can look at the dollar beta allocated to the different areas relative to the surplus, and translate that to the risk capital. Then we can compare that to the equity in the firm versus the return on risk-adjusted capital versus the return on surplus.

There are a lot of applications of VAR. I think it's in its infancy in terms of applications. Getting into the game and looking at the various investments is a good starting point. Spreading out and starting to consider your whole liability profile becomes a lot more complicated, but we're doing it with several firms already. It's giving them some measure of controls that they need to implement in terms of the different functions within the organization. I think that's providing a lot of utility in terms of the discussions of what they need to do to better manage risks.

As we go through this process more and more, I'm hoping that there will be more publicly available information. Perhaps our company will be able to provide that more readily. But that will address the generic needs of different companies and

institutions that can be used as a starting point for peer group comparison, and further tweaking to data that's available to private entities.

We've already been doing that sort of work in great detail with some large life companies, banks, and Treasury units where they're looking at risk and performance. Our colleagues over at Rogers Casey are doing this for plan sponsors already.

I think the ultimate goal is to look at increasing the value of the firm through improved investments, better A/L management strategy, and optimal capital and asset allocation. And I think VAR can certainly help in that area.

From the Floor: I have one question. The VAR really came from the banking field and is getting into the insurance field. It seems to be mostly focused on the asset side. Correct me if I'm wrong. How sophisticated is the liability modeling in the bank use of VAR? If it's going to tag onto this, how sophisticated does the liability modeling need to be? It seems to me that with the Monte Carlo method the people have been weaker with the liability modeling.

Mr. Abbott: There's no question that the starting area for applying VAR was at the trading desk level where they were looking at their short-term exposure. It was not used, and has not yet been used much at the treasury level of the bank. It has not been built up to the asset/liability management committee. In all but a few institutions, I think that this measure is even new to them. You're absolutely right. When you're dealing with very complex prepayment models, or flow of funds models in the banking institution, they haven't gotten all that coordinated yet. We've worked with several to produce results, and we've simplified the models. We're not as proficient with accounting issues. These details really complicate and slow down things. That's the same thing for the insurance side. When you deal with the models that you're using for regulatory reporting, you're going to get bogged down in terms of the details for this information. But you're not going to get control over parametric sensitivities.

So we built special models independent of those details to get at the risk drivers. Our first approximation is achieved by working very closely with the actuaries in those organizations so that we can get basic numbers out in front of people, and then refine those models. It's essentially a feedback loop. We're getting information to those companies, and I hope that our successes there will expand to other organizations.

Ms. Orlando: One thing I just might add is that it's very important when you're looking at insurance companies and trying to, first, come up with a framework for doing this on both the A/L side, and then trying to implement something. Now you

look at your assumptions and make sure that they are consistent among your various business units and among the other people who are going to be supplying numbers to these rolled up results. We found a lot of cases where there are certain assumptions used for cash-flow testing, and those are fine. Sometimes companies even have a hard time getting those assumptions consistent.

You have to make sure that you go through the assumptions that are being used for this kind of a calculation and you must make sure that they're consistent, because we find in each business unit that there are people with proprietary models, and everybody's calculating his or her own thing. Everybody loves his or her model and hates everybody else's model. You have to make sure that they're all basically putting in the same input.

Mr. Abbott: I think that's a very good point. You need to make sure that you're dealing with apples and apples to the best extent that you can. That's another good reason to roll things up to one model where you know that the risk drivers are all interacting correctly. Developing the various covariance matrices is extremely difficult because there's sparse or no data in a lot of cases. Another thing is that we've actually worked with clients that don't apply only the covariance matrices we generated from the historical data. We use a GARCH methodology that weights more toward the present. They'd like to see what happened over certain events, so we calculate a covariance matrix over the market crash of 1987, or over the peso evaluation or the recent Asian debacle. Getting some realistic sense of how that covariance matrix has changed from period to period also gives them the sense of what their losses could be for some actual event.

From the Floor: I think it was mentioned a couple of times. The VAR works better over a short-term horizon rather than a long-term horizon. How short is short?

Ms. Orlando: Well, a typical VAR is calculated over a one-day to a ten-day holding period. That's short. That's too short for insurance companies. You can probably get some decent results over a six-month or a one-year period, but I wouldn't do it on a VAR basis. I would do it on an earnings-at-risk or a surplus-at-risk basis. That's much more meaningful to an insurance company. Otherwise, it's just a number for them. It really doesn't tell them anything.

From the Floor: Speaking of horizons, how would you reflect the value of new business coming into calculations, realizing that businesses have a 12% growth rate over a five-year period? Not only assets, but premiums written can usually double.

Mr. Abbott: Yes. The corporate model should consider growth.

From the Floor: I would think they would want to play a role in the use of the system and also the horizon they're looking at.

Mr. Abbott: The models that we use actually can add either a run-off view a growth view to the liabilities. Obviously, if you're doing that on the liability side, you need to have corresponding future purchases of assets as they would be in these particular cases or various scenarios. That becomes much more complex. It slows down the model. You're dealing with a multiperiod simulation essentially, but it certainly can provide a lot more value. What we like to do is take a look at the exposures that one would have with that growth in business and apply that more generically to the asset classes and future purchases rather than deal with specific new CMOs and other things. We'll just deal with a general category, so we'll create a sensitive index to the risk drivers at those futures points in time.

From the Floor: Does the assumption about the growth of new business affect the usability of a variance/covariance model over the Monte Carlo approach?

Mr. Abbott: That's a good question. I haven't thought about that. Certainly, you were doing Monte Carlo simulations typically to model the liabilities, so I'll have to look at that a little bit more. We're dealing with the parametric model that's now looking at all of these different path-dependent liabilities or assets in the simulations. We don't have to deal with other assets that are far more static in that fashion. We can just deal with them parametrically and get at the results. It's a lot faster to deal with those as one set, and the path dependent ones as a separate entity, and then net them together.

But certainly, doing things through the pure Monte Carlo approach can give you some better sensitivity when you do have to reinvest or disinvest new premiums being sold. Persistency is necessary to value and address the profitability of the ongoing business.

CHART 1
CORRELATIONS BETWEEN US STOCKS AND TREASURY BONDS

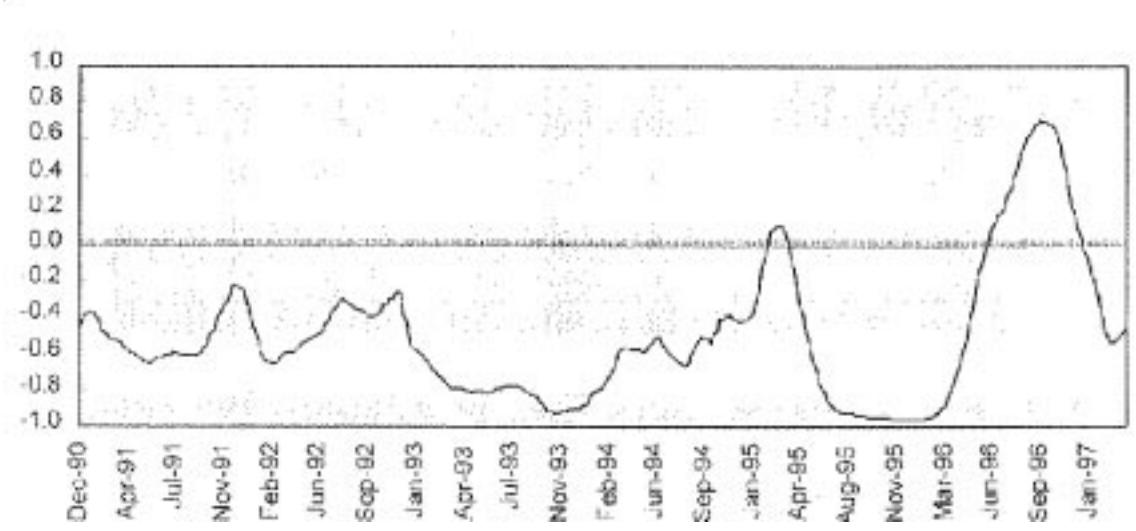

CHART 2
VOLATILITY OF S&P INDEX PRICES

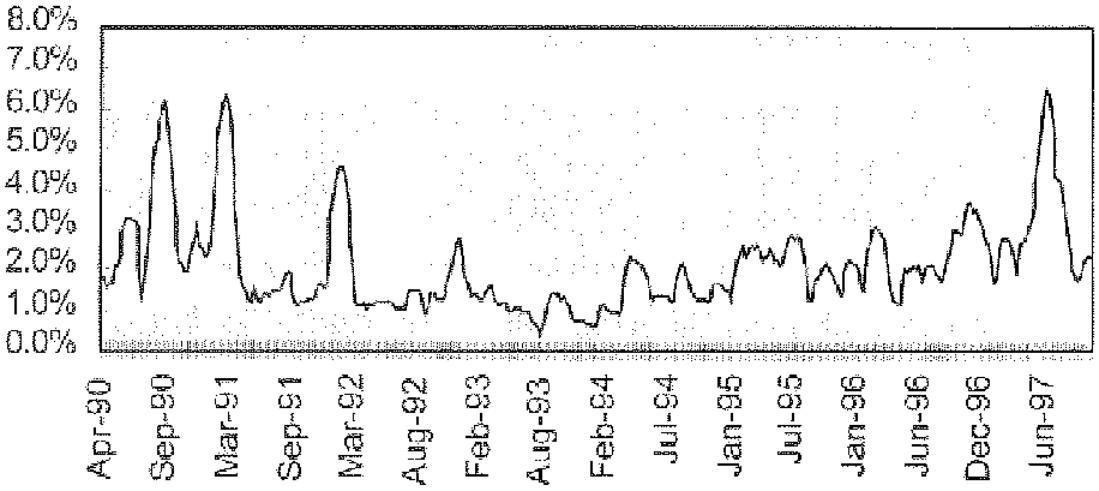

CHART 3
SAMPLE VALUE-AT-RISK

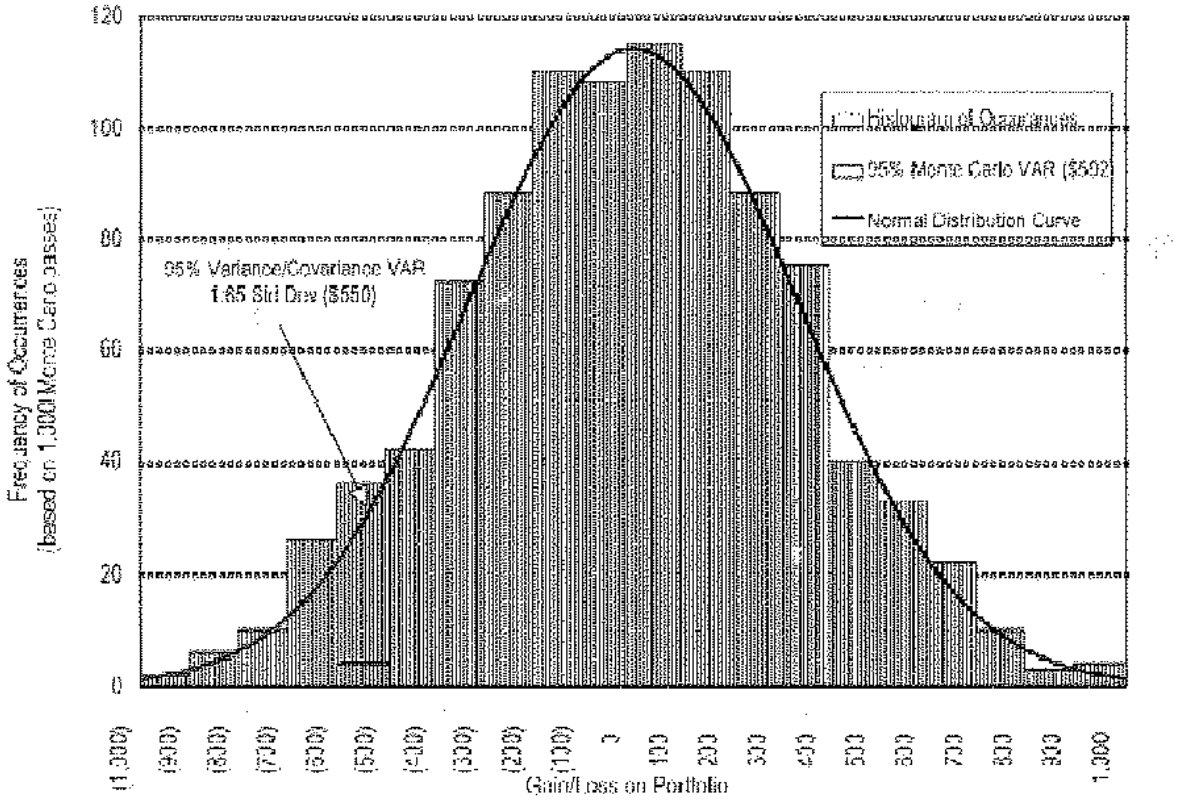

CHART 4
KEY RATE DURATION SHIFTS

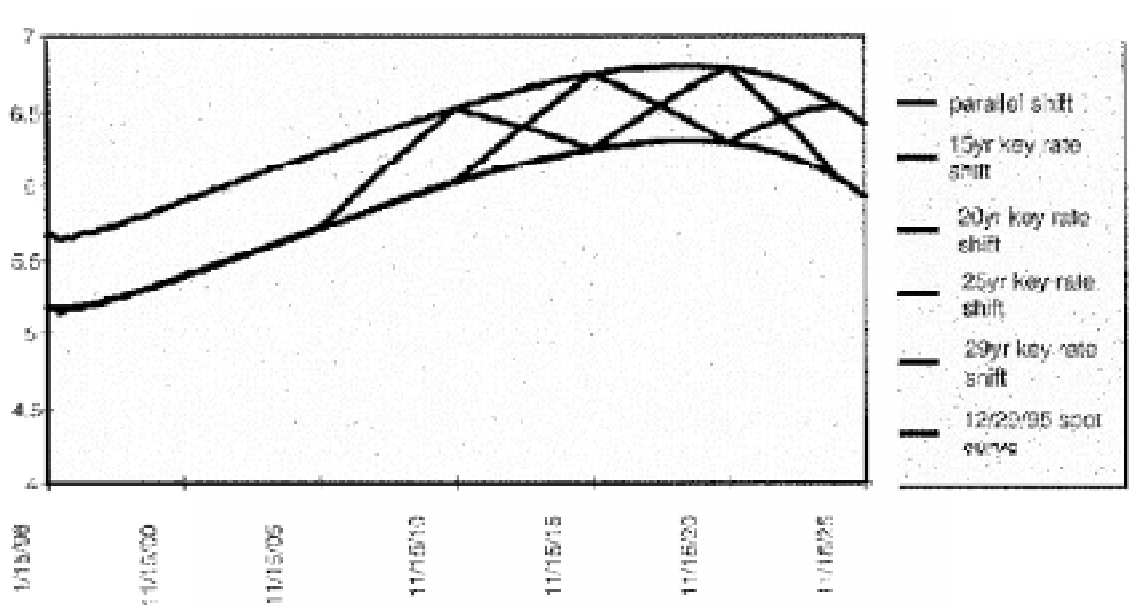

CHART 5
EY RATE DURATIONS
9% VS. 8% 30-YEAR CALLABLE CORPORATE BONDS