

**RECORD OF SOCIETY OF ACTUARIES  
1995 VOL. 21 NO. 4B**

**DEVELOPING “DOLLARS-AT-RISK” LIMITS**

Moderator: CINDY L. FORBES  
Panelists: RICHARD W. HARRIS  
              GLYN A. HOLTON  
Recorder: FIDELIA CHEN

*The banks are moving toward “dollars-at-risk” systems for risk management. Does it have applicability to the insurance industry?*

MR. GLYN A. HOLTON: I was just thinking back to the 1980s when I used to work at Metropolitan Life. I worked in the group annuity area. We were pricing guaranteed investment contracts (GICs) and other group annuities. Back in those days, the pricing was really tight. You’d do a competitive bid and maybe three, maybe four, or maybe ten of them were carriers. You’d usually see bids within five basis points of each other, and occasionally, there were the outliers. One company would come in and beat you by 20 basis points. It was very frustrating. It would be very easy to say that it left a lot of margin on the table, but at the same time you wondered what it was doing. How did it beat you like that? It was a disturbing question, and there weren’t any simple answers to it.

Those were two comparatives that we really worried about back in those days. One was the banks. We didn’t have any idea what they were doing regarding bank investment contracts (BICs). Of course, we worried about Executive Life as well. We never did beat Executive Life. We eventually just stopped bidding against it, but it was always great to take an opportunity, if you could, and sit down with a banker and find out what was on his or her mind. How was the bank pricing these things? Not that its way was better or worse than ours, but it gave us some insight into what it was doing.

With BICs in particular, I learned that bankers were pricing them on the trading floor, and we’d be actually pricing these things in the office of a senior vice president. When bankers would close a BIC, they would immediately book a profit. When we would close a GIC, we would immediately book surplus strain. So they were quite different approaches. That was back in the 1980s. I went on and worked in investment management, I worked in banking, and now I’m back consulting to the life industry.

What I really enjoy is the opportunity to talk about how other people think about these things. How do other industries respond to some of the issues that we face? That’s what we’re going to be talking about—value at risk, also sometimes called dollars at risk. It’s a concept that is really sweeping Wall Street.

Why are people concerned about value at risk? Why is this a popular new risk measure? Wall Street financial institutions have always been about managing risk! What has changed? Why the focus today?

What’s changing is how we manage risk. In the past, we’ve always managed risk on a local level. A trader would manage the risk in his or her portfolio through a portfolio manager. An actuary would manage the risk in a particular book of business. I think

the world is changing, and the evidence that we see coming across our desks in *The Wall Street Journal* everyday are some of the spectacular losses that organizations are having—Orange County, the Barings Brothers, and the Daiwa Bank.

People are realizing that our traditional approach in managing risk no longer works. You can't manage risk on a tactical level. You need to manage it on a global level and across products. You need to do it in a centralized way. There's still the role for the actuaries, the portfolio managers, and the traders managing risk on a tactical level. We need to roll that risk up and manage it on a global level. To do that we're going to need a new way of measuring risk.

Let's take a look at some of the risk measures that are used. I have three general classifications of risk measure: factor sensitivities, single-scenario analysis, and statistical risk measures.

We're familiar with factor sensitivities in the form of duration and convexity, which measure your sensitivity to a particular risk factor—that risk factor being the parallel shift of the yield curve. Anyone who's familiar with option pricing theory is familiar with delta, gamma, beta, and rho, which measure sensitivity to particular risk factors, first-order or second-order sensitivity to movements in the underlier, sensitivity to movements in implied volatilities, etc. Those are risk factors.

Risk analysis techniques, or what I call single-scenario risk measures include cash-flow testing. Dynamic solvency testing (DST) also falls into this category. You specify a specific scenario that impacts a variety of risk factors, possibly over an extended horizon, and you look at how your cash flows will evolve, how your surplus will evolve over time if this particular scenario were realized. Of course, you can go on and do several scenarios if you want, but you analyze one scenario at a time.

The idea behind statistical risk measures is that you're concerned about the market risk in your portfolio, maybe over a short horizon (a day) or maybe over a longer horizon (a year), and you characterize the probability distributions for the market value of your portfolio at a future horizon. Then you characterize that entire probability distribution, with maybe one or two statistics. I think the most obvious example of a statistical risk measure is historical volatility. Historical volatility is just a standard deviation of that probability distribution from the future market value of your portfolio. The value at risk is also a statistical risk measure.

I'd like to point out two examples of the limitations of existing risk measures that can be used. Let's start off with an example of duration/convexity, which is commonly used by insurance companies for measuring risk. We're going to consider two different portfolios. The first portfolio is very simple. It consists of a single five-year zero-coupon Treasury bond. Because it's a zero-coupon bond, its duration is precisely five years. It also has convexity; in this case, it is 34.8.

We construct another portfolio. This is a cash portfolio. Typically, a cash portfolio will have a very short duration. If the average maturity in the portfolio is three months, its duration will be about three months. We've added a floor to this cash portfolio. As interest rates decline, the value of that floor increases. What that means is that the overall portfolio will tend to act somewhat like the five-year zero-coupon

## DEVELOPING “DOLLARS AT RISK” LIMITS

bond as interest rates rise or decline. We can construct this portfolio in such a manner that the cash combined with the floor will have a duration, option adjusted, of exactly five years, and it will have convexity of exactly 34.8 years. My question to you is which portfolio is more risky? Is one portfolio more risky than the other? Or are the two portfolios equally risky?

From a duration/convexity standpoint, we might conclude that they were equally risky. In fact, they have completely different risk characteristics. The reason is that the five-year zero coupon bond has a duration of five years and a convexity of 34.8 years, but that's because of exposure to the five-year interest rate. The second portfolio has a five-year duration and convexity of 34.8, but that's because of exposure to the three-month London InterBank Offered Rate (LIBOR). The three-month LIBOR is substantially more volatile than the five-year Treasury rate. Often you'll find it to be twice as volatile. That means that your second portfolio has twice the risk of the first portfolio, but your duration convexity fails to capture it. Why? Because it's not a statistical risk measure; it's a factor sensitivity.

Let's take another example. Let's look at historical volatility as a risk measure. The robust risk measure captures all sorts of different sources of risk by summarizing the impact on your market value irrespective of the source of risk. Suppose you've been tracking a portfolio and computed its historical volatility at 20%. Maybe we've looked at price fluctuations over the last 100 trading days and the standard deviation of 20%. We know how risky that portfolio is.

But suppose we rebalance it? We're going to reduce our exposure to domestic equities and start investing overseas. Investing overseas adds diversification to the portfolio, which reduces the risk. It also adds currency risk, which increases the risk. What is the riskiness of the portfolio now? We don't know. Historical volatility can't answer that question for us.

The problem with historical volatility is it is a retrospective risk measure. It is not a prospective risk measure. It's a powerful risk measure, but it will not always address your problem. If you're a trader who is trading real time, you need a real-time risk measure. You want to know how much risk you have right now. If you're hedging an option position and you've blown the hedge, you want to know now so you can fix the problem.

Let's talk about value at risk. Value at risk is a statistical risk measure. It's designed to capture all the factor sensitivities that impact your portfolio. It's also designed to capture factor volatilities and all hedging and diversification effects. In addition, value at risk can be updated if the portfolio is traded.

What's the definition of value at risk? Value at risk is the amount of money such that the current portfolio would not be expected to lose more than that amount of money 95 days out of 100. It's a very general, simple, and powerful definition.

Value at risk is a statistical risk measure because it's in the upper bound of the confidence interval. You think about the probability distribution for the future price of your portfolio maybe a day, a week, or a month from now, and you give the upper bound a 95% confidence interval. People use different confidence intervals—90%,

95%, or 99%—and different time intervals—looking at the value at risk over one day, one week, six months, and so forth.

Value at risk captures all sources of market risk. It can be used across products. You can apply this to a portfolio that perhaps contains equities, fixed income, real estate, a Friday night poker game—there are no limitations on it whatsoever. The question is how to go about computing it.

Before I talk about the calculation, I'd like to talk quickly about the advantages. One is, as I just mentioned, the robustness—it summarizes broad-spectrum risk with a single number. The second is the ability to update your measure as the portfolio is traded. The third is it gives you very intuitive and meaningful output—you present to senior management or the board of directors and you do not have to explain what gamma and deltas are.

There are a few methodologies for calculating value at risk. However, I would like to talk to you about closed-form value at risk, which is a simple calculation widely used in banking today. The one significant shortcoming of it is it cannot be applied to portfolios that have significant gamma or convexity risk, and that's a significant shortcoming. I'll show you how it works.

The first step in doing a value-at-risk calculation is to specify key variables. These key variables are the sources of risk in your portfolio. For example, if you had a currency portfolio, the key variable would be the exchange rates on all the significant currencies you were trading. If you were trading currencies forward, you'd also be taking yield curve risk—domestic yield curve risk as well as foreign yield curve risk. So a complete set of key variables might be the various exchange rates that you're trading: foreign and domestic interest rates, and perhaps a one-month LIBOR and a one-year LIBOR for each exchange rate. If you have options in your portfolio, you have to think about implied volatilities and other such things. You want to get a complete set of key variables such that you can value your entire portfolio, by using those key variables.

Step two is to make two unique fundamental assumptions for closed-form value at risk; they're not necessary for value at risk in general. The first assumption is that your key variables are jointly normally distributed. Some people may be uncomfortable with this assumption. You might think some key variables would be more appropriately modeled with the lognormal distributions. It's not a very big assumption, though. It's a fairly reasonable assumption for risk measurement purposes. You might not want to price or look at a swap by using the normal assumption, but if you want to measure your value at risk to within 0.5%, the normal assumption should be sufficient.

The second assumption is much more significant in that the value of your portfolio is linearly dependent on a key variable. This assumption is going to enable us to come to a very simple formula for value at risk. At the same time, it's going to make that formula meaningless if you have gamma or convexity in your portfolio; so this is a critical assumption.

## DEVELOPING “DOLLARS AT RISK” LIMITS

Step three is to characterize the probability distribution for our key variables. We have assumed that they are jointly normally distributed; therefore, characterizing the distribution is fairly easy. We need to specify a mean, a standard deviation, and a correlation to each of the key variables.

In a banking environment you'd calculate value at risk over a day. If you want, you can assign a mean that would be equal to the risk-free return plus a spread for any systematic risk over a day. Typically, however, people assume it's the expected return thereof; it's a simplifying assumption but it's not necessary, though.

You have to specify standard deviation correlations for each key variable. J.P. Morgan's *Risk Matrix Data Base* publishes standard deviations and correlations for a wide variety of key variables today. Many organizations are starting to use these to support their value-at-risk calculations by getting on the Internet.

With these variables, we have fully characterized the joint and normal distributions for our key variables. This information is half the solution; it fully characterizes the markets and the riskiness that exists in the markets.

Step four is to determine the portfolio's sensitivities to each one of the key variables, the second half of the solution. You need to find the first derivative for the portfolio's value taken with respect to each key variable. Because derivatives add linearly, that's the same as taking the first derivative of each contract value with respect to each key variable and then summing across all the contracts in the portfolio.

This calculation is not an approximation for the sensitivities. It's the precise formula for the sensitivities because we have assumed a linear relationship, so the second derivatives would all be zero. In step three, we fully described the markets, the key variables that impact the portfolio, and now we've described the portfolio's sensitivities to those markets.

$$S_i = dP/dV_i = \sum_j dC_j/dV_i$$

where  $P$  is the portfolio's value and  $C_j$  is a contract value for contract  $j$  in the portfolio.

In step five, we can conclude that the price of our portfolio is going to be normally distributed because we assume that our key variables were jointly normally distributed, and we assume that our portfolio's value was linearly dependent on those key variables. Linear relationships preserve normal distributions and therefore, the portfolio will have a normal distribution as well. That normal distribution will have a mean of zero because we assume that this key variable had a mean of zero, and linear relationships will preserve mean. I'm using the mathematician's definition of linear relationships.


Finally, we have a formula for the standard deviation of our normal distribution of the portfolio. There's nothing new about this formula. It's just a standard formula for summing standard deviations. We require no assumptions to apply this formula. We pull it straight out of our standard statistics text. We're adding the various volatilities

in correlations of the key variables weighted for the sensitivities. We get a precise number for the standard deviation of the price of our portfolio. We have fully characterized the probability distribution for the price of our portfolio. Standard deviation given by:

$$\sigma_p = \sqrt{\sum_i (S_i \sigma_i)^2 - 2 \sum_{i < j} S_i \sigma_i S_j \sigma_j \rho_{ij}}$$

In step six (Chart 1) we've shown an illustration of a normal distribution, and our value at risk is just the upper bound. The upper bound in that confidence interval is just a multiple of your standard deviation, if it's a normal distribution. For a 90% confidence interval it would be 1.39 times your standard deviation. If you're computing value at risk at the 95% confidence level, I think it's 1.65 times your standard deviation. But your problem is done. We've found a solution for value at risk.

CHART 1  
PORTFOLIO LOSS/GAIN



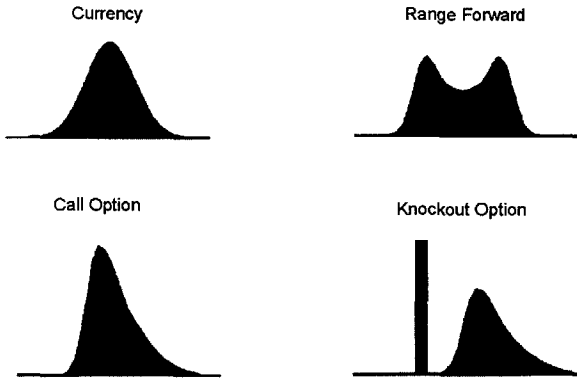
In step seven, we can make the calculation real time if required. The individual standard deviations and correlations for the key variables remain constant throughout the day. J.P. Morgan posts them based on historical valuations each night, so those will not change during the day. What changes during the day is your sensitivities, the S of I. As you trade a portfolio, you can update those sensitivities and just run them through the formula. You can calculate value at risk for a reasonable portfolio on a 486-based PC on a real-time basis. That's the power of closed-form value at risk. Make it REAL-TIME:

$$\sigma_p = \sqrt{\sum_i (S_i \sigma_i)^2 - 2 \sum_{i < j} S_i \sigma_i S_j \sigma_j \rho_{ij}}$$

However, closed-form value at risk cannot be applied in all situations. Recall our normality assumption, our linearity assumption. If you have convexity or gamma in your portfolio, you cannot apply closed-form value at risk. The notion of value at risk is still meaningful. It is still a powerful measure, but you're going to have to find a new way of measuring it. Four examples are shown in Chart 2.

## DEVELOPING “DOLLARS AT RISK” LIMITS

### CHART 2 DISADVANTAGE OF CLOSED FORM VAR: NOT ALL PORTFOLIOS SATISFY THE SIMPLIFYING ASSUMPTIONS



In the upper left-hand corner, we're looking at a currency whose value we can reasonably treat as being normally distributed. You could have currencies in your portfolio and that, perhaps, could be all you have. Then closed-form value at risk would be a reasonable solution.

In the lower left-hand corner, though, is the payout function, the probability distribution for price, a call option, which is a skewed function. This call has no downside risk. If you apply the closed-form value-at-risk model to that call option, it would fit a normal distribution to that curve, and it would say that you do have downside risk when, in fact, you don't. People buy options because they don't have downside risk; you don't want your risk management software suggesting that they do have downside risk. You would get a misestimation of what your exposure is.

In the upper right-hand corner is a range forward. There's a probability distribution that's dramatically nonnormal. A closed-form value-at-risk model will not work for this either. In the lower right-hand corner is a knockout option, which is dramatically nonnormal. You can apply value at risk to these instruments, but you'll need to use a different model.

For such instruments, the best solution is typically a Monte-Carlo-type value-at-risk analysis. I'm not going to walk you through that methodology, but it is based very much on the same sort of analysis as the closed-form value at risk. You start off with key variables and then you perform a Monte Carlo simulation. You can perform a Monte Carlo simulation on a portfolio, with effectively no simplifying assumptions, and get a very accurate measure of what your value at risk is. Monte Carlo simulation requires no joint-normal or linearity assumptions. It can handle nonlinear instruments, such as options, mortgage-backed securities, structured notes, or so forth. There are ways to speed up Monte Carlo simulations, particularly it's called reflex technology, so that you can capture most of the calculations, run them overnight, and then just update them during the day. You can actually do reflex Monte Carlo simulation on a real-time basis. It takes a lot of work, though.

To summarize, value at risk is a robust risk measure. It's applicable across products, across all different types of sources of risk. It has an intuitively appealing output, and it is meaningful to nontechnical people. It also captures all hedging and diversification effects because you're taking into account those correlations. In many cases, it can be performed on a real-time basis.

I've written a paper, *Closed Form Value At Risk*, which, to a large extent, parallels the discussion I've had of how closed-form models work. It does take a slightly different approach to the material than what I've taken right now. I've tried to make this presentation intuitively appealing. The paper is technically rigorous and takes a slightly more general approach to the methodology. If you read that paper and find it useful, I'm writing two other papers—one on Monte Carlo value at risk and another on reflex value at risk—and they'll be coming out in probably the next two to three months. If you're interested, I'd be happy to pass them on to you.

MS. CINDY L. FORBES: Now I'm going to quickly go over how some of these applications or concepts might be applied to a life insurance company. Clearly, insurance companies that operate multinationally also have a need to add up their various sources of risk across all their operations. From that point of view, you might have portfolios around the world, for which you would want to know your overall risk position; and that's really what brought banks into the business of developing value at risk. Another idea for insurance companies, though, is that they can look at their entire balance sheets and summarize the other risk factors that are on the balance sheet, such as credit, mortality, and morbidity.

You can create a communication tool for senior management that tells management what its value at risk is for the entire corporate entity and that also decomposes it in terms of telling what its biggest risks are on the balance sheet. That would enable management to determine whether they believe that they are getting appropriate returns from the risks that they have and also enable them to take risks with what they believe are their core competencies. It allows senior management to take a look at how much earnings volatility they're willing to live with and whether current practices fall within that range. Are they taking too much risk or too little risk?

I think value at risk is a natural fit for actuaries and for the insurance industry as a whole, given the technical nature of the business and the profession. It's an easy knowledge extension from the banking world to the life insurance model. Perhaps the most ambitious and exciting possible application is coming up with a total communication tool for senior management that tells them their risk position, but there many other ways that value at risk can be used in an insurance environment.

For instance, we have just used the same methodology to determine portfolio trading limits for actively traded portfolios. Rich will tell you a little bit more about how we did that. It can also be used to validate or develop your own internal risk-based capital requirements, based on what you believe the risk position or risk characteristics of each asset/liability are.

However, no matter what the similarities and the applications, there are some key differences between insurance companies and banks. One thing is that banks typically calculate value at risk over relatively short periods of time. Banks tend to think of



## DEVELOPING “DOLLARS AT RISK” LIMITS

value at risk as being the amount that they would have at risk over a period of time, and that period of time is how long it would take them to dispose of that asset without having any noticeable market impact. That means that typically they will look at risk measures or value at risk over one day or 30 days. That would be a typical time frame, and many insurance companies might think that that's not the right time frame. Then you need to look at different methods of calculating value at risk that apply to longer holding periods.

The other thing that is different with insurance companies, and it relates to the holding period as well, is that insurance companies typically have many more illiquid assets in their asset portfolio. What does that mean? We all know that if we have an interest rate risk position, we could fairly quickly go out in the market and overlay it to put us back to risk-neutral, if that's what we want to do. So from an interest rate risk point of view, our time frame is not all that different perhaps from the banking industry.

However, when you get to some of your particular assets, you may think that you can't dispose of them or wouldn't want to dispose of them. So what is the appropriate time frame? You could securitize most assets these days in six months, if you wanted to, but maybe that's not the appropriate way of looking at the question.

Value at risk is really a way of communicating to management how much of this quarter's earnings or this year's earnings management has at risk; and that's what management tends to think of. What could be the possible impact on my earnings this quarter or this year? You might argue, regardless of your holding period, that you wouldn't really want to measure your value at risk over a longer period than what management thought of in terms of managing its earnings.

Another key difference between banks and insurance companies is that banks tend to apply value at risk to their entire balance sheet. Each component part is measured separately. Insurance companies tend to match up assets and liabilities, put them into a separate segment or portfolios, and tend to look at the net position—assets minus liabilities. Therefore, the model developed for the insurance industry will have to be applicable to the liability side of the balance sheet as well and be able to show you the impact of any scenario analysis that you're doing on the liability side as well as on the asset side at the same time. Clearly, many liabilities don't fall into a closed-form value-at-risk calculation; they're very complex and certainly aren't linear. You're going to have to look at Monte Carlo-type techniques to be able to model them.

The Investment Section Council has just authorized a research proposal on how to extend value at risk to the life insurance industry. We're just getting started, but I'm hopeful that we'll be able to come up with a paper that shows people how they can best go from the banking model to the insurance model. Now I'd like to turn it over to Rich to give his overview of how we applied dollars at risk at Manulife.

MR. RICHARD W. HARRIS: You've heard Glyn speak about how banks have applied the dollars-at-risk concept over a short horizon in a closed-form type of solution. You've also heard Cindy compare the applicability of dollars at risk in the insurance company environment to its applicability in the bank environment. I want to outline a practical example of how we have started to develop dollars at risk at our

company and how we plan to integrate the concept into our management and decision-making process.

I will discuss the management framework that we have designed around dollars at risk and then talk a bit about the mathematical framework which, for the insurance industry, centers around creating probability distributions stochastically rather than using theoretical closed-form distributions as in most bank applications. I will outline some of the other issues around dollars at risk that we considered in our development, but which we don't plan to implement in our initial stages. Then, I will finish with some practical observations on implementing a dollars-at-risk system in the insurance company environment.

Let me start out by giving some background on our specific dollars-at-risk application. We chose to look at our new money funds in both Canada and the U.S., consisting mainly of 401(k) and GIC-type products, and a large block of payout annuities—all backed by bond and commercial mortgage portfolios. Risk management is currently being performed on this block by using the traditional duration and convexity measures on a deterministic basis without tying the risk position back to any type of probability loss and without setting any dollar limits on the amount of risk that can be taken.

We found we had a large proportion of the company's assets backing reasonably simple products, with a potential for significant market risk exposure under any kind of active trading. This is the direct application of limiting market risk that Cindy mentioned. We considered application of dollars at risk to several other different risks—currency, mortality, credit—but we decided to concentrate on market risk—both interest rate and sector spread risk, and I'll limit my discussion to the interest rate risk on this portfolio.

I will use the term *risk limit* the same way Glyn used the term *value at risk*—the dollar point on the distribution or the upper bound on the confidence interval that is chosen as the maximum level within which risk may be taken. I will refer to the term *dollars at risk* as the dollar measure of the current position of the portfolio, the current market risk, including any year-to-date actual losses, if applicable. Both terms will be used interchangeably. I would like to talk first about developing the management framework around dollars at risk before getting into the mathematical development. That's actually how we approached the problem at our company. In deciding how dollars at risk was going to work, we thought that the main challenge was to design a management framework that tied together all the various risk parameters so that we could rework all the pieces of the puzzle and put them in place. We debated long and hard about all the various pieces and how they should work until a final product was obtained that was consistent and workable.

The first issue we had to deal with was the actual dollar measure that we were going to use. What were the dollars at risk that we were talking about? We considered statutory earnings, economic value or present value of future profit, and also dollars of capital. In the end, we wanted a measure that considered the impact of both sides of the balance sheet and a measure that was highly visible, at least in the final product, and easily understood by senior management. Based on that, we chose statutory earnings at risk as our dollar measure.

## DEVELOPING “DOLLARS AT RISK” LIMITS

We then considered the time frame over which a management cycle would extend. We considered an open time frame in which the dollar limit would be an available risk exposure no matter what the current conditions or past losses were. We also considered different set time frames, such as monthly, quarterly, and annually, which, as Cindy mentioned before, are a lot longer than what a bank would normally consider. A combination of a short-term limit based on a dollar measure and a longer-term limit based on another was also considered. In the end, we finally decided to use dollars of annual earnings at risk, which tie into the business planning cycle and our earnings management approach.

We also debated long and hard about whether the limit should be open in terms of consideration of past gains or losses or whether the limit should decrease over the year if actual losses were sustained. This is the question of, if your limit is \$10 million and you make \$10 million on an interest rate, is your limit now \$20 million for the year? Conversely, if you lose \$5 million, is the remaining limit \$10 million, or is the remaining limit \$5 million?

We chose the latter approach, what we call a closed approach, in which actual losses reduce the limit and actual gains leave the limit unchanged. We put, in essence, an upper bound on the amount of negative earnings volatility we have within the probability distribution that we generated.

We also spent time designing the reporting and monitoring process that would be used for risk limits. We thought that the impact of trading activity should be monitored on a daily basis at the desk level, and that formal reporting should be made to asset/liability coordination (ALCO) and senior management at least on a monthly basis.

The real value we saw in risk limits as a management tool was the ability to stimulate communication and risk management action based on certain trigger point levels of current dollars at risk in the current mismatched position compared with the total dollar risk limit that we had set. In addition, rather than being a purely technical exercise, we left room for flexibility to change current limits based on management's appetite for risk due to various internal or external factors into the design of the process.

The above was achieved by designing an annual proposal process to tie in with the planning cycle. The flexibility for certain levels of management to be able to change the ultimate limit or to reallocate risk dollars among certain risks or funds based at certain points of time during the year, based on their current appetite for risk, was built in.

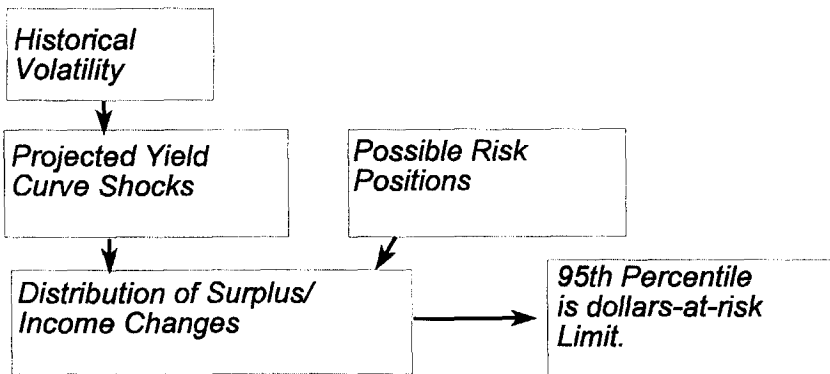
In summary, we designed a framework for managing dollars of annual earnings at risk to control negative earnings volatility from market risk in our new money funds. We also developed a trigger point system that would stimulate management discussion and risk management action as market dollars at risk levels approached the ultimate risk limit.

Next, I briefly want to describe the mathematical framework on which we built our dollars-at-risk process. As I alluded to earlier, the mathematical process was one of

creating a probability distribution of outcomes to set a confidence interval within which risk limits would operate. There are four key building blocks of this mathematical process. First is the size of the risk positions necessary to flexibly manage the business. These positions were designed and stated in terms of the mismatch between assets and liabilities, so that we're looking at the net earnings position. Second is the historical volatility of, in this case, the interest rate risk. Third is the ability to take the historical data and stochastically generate a set of interest rate shocks which, when applied to a specific interest rate risk position, create a distribution of earnings, gains or losses. I won't go into a lot of detail around our stochastic interest rate generation process, but what we did was base our model on James Tilley's paper, "An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators" (TSA XLIV,1993). That, in general terms, takes a lognormal, mean-reverting process and models out the interest rate paths to the horizon that you choose. The fourth and final element is the level of risk usually expressed in terms of confidence intervals or a number of standard deviations within the distribution—Glyn mentioned 90%, 95%, 99% confidence intervals—which need to be decided upon to determine the ultimate risk limit level.

Chart 3 shows graphically the theoretical process of using historical volatility to generate interest rate shocks or projected yield curves and applying these projected yield curve positions to generate a distribution of income changes. We applied the shocks to five different interest rate risk positions that the bond desk, in our case the interest rate risk managers, thought they might take under certain situations. We chose the 95% confidence level on the most adverse distribution as our risk limit level. Applying this same confidence interval to the distribution generator from our current interest rate risk position or our current portfolio of asset/liability cash flows and adding any year-to-date losses sustained gives the current dollars at risk, which can then be compared with the ultimate risk limit level. In other words, the mathematical framework lends itself directly to calculating both an overall risk limit level and current dollars at risk just by substituting the risk position or the mismatch between assets and liabilities.

CHART 3  
INTEREST RATE RISK  
DOLLARS-AT-RISK LIMIT CALCULATION



## DEVELOPING “DOLLARS AT RISK” LIMITS

Some of the issues that we discussed, but decided to leave out of our initial application were:

1. Correlation between risks and between asset types. This is an area where good historical data and the flexibility to adjust to changes in the mathematical relationships are key. We decided initially that we would take a more conservative and easier approach to our risk limit implementation—in leaving out correlations and basically adding the risk factors together between the portfolios and between the different asset types.
2. Other key risks, such as credit, currency risk, etc., were considered, but we decided to leave them out because we wanted to set up a framework and implement a simple application initially with the big bang market risk in which we had a lot of vulnerability in this portfolio and not deal with too much complexity initially. It should be emphasized, though, that the framework that we have set up can easily be extended to these other risks and other types of assets.
3. The transfer of risk dollars was not built in. You can set up a complicated system of transferring risk dollars between portfolios, between funds, to take advantage of certain economic events, or a certain appetite for dollars at risk, or for market timing or various timing incidents. We decided that building this in was a complication that we didn't want to introduce initially.

Finally, I would like to end with some practical observations. Dollars at risk is a very useful tool for quantifying and managing risks in dollar terms that can be combined across funds, across risks, etc. It's also a tool that is easily understood and can be communicated in terms of probability levels without all the complex mathematical terms.

The theoretical development of dollars at risk lends itself very readily to a *communication and risk management framework built around the risk limit that is agreed to and set*. The comparison of current dollars at risk to an overall risk limit gives immediate feedback in probabilistic terms as to the amount of risk the company is exposed to. The theoretical development is also generic, so it can be easily expanded to other risks and other funds.

The four building blocks of risk position, volatility, distribution of results, and probabilistic level of risk are common to all dollars-at-risk applications. Again, we created the probability distribution in this case. We didn't rely on closed-form assumptions, so when expanding to other risks we will look at the risks available and the probabilities thereof and develop a new distribution of results.

As well, there's a practical reality. It took us eight months to develop this simple application. I know there are broader applications going on in other financial institutions that have been in the order of two to three years to fully implement. We do think, however, that the groundwork we've laid for the application and eventual extension of the dollars-at-risk concept was well worth it.

MR. HARRY H. PANJER: I didn't know anything about the dollars-at-risk concept until I talked to Cindy for a few minutes the other night, so I came here to learn

something. I will just share my observations. The dollars-at-risk concept seems to be what actuaries have been doing all along in what has been consequently called risk theory; that is, to try to look at the distribution of losses from whatever source, construct an appropriate model, and look at the amount of capital required to support those losses.

When you look at dollars at risk as a probabilistic point of view, as presented here, that's simply looking at the 95th percentile of that distribution. The dollars simply represent the amount of capital that you can potentially lose in that time horizon. So, fundamentally, that is no different from what actuaries have done. The question then becomes one of constructing the model.

The first model that Mr. Holton presented is essentially a vastly oversimplified model, because it's simply a standard multivariate normal model that all actuarial students learn in their third year at the university, and a linear combination of elements in a multivariate normal model. He claimed that the model is robust; well, the 95th percentile is robust.

Actually, I would suggest that *robust* is not the right term. It is exact in that context, in the context of the model assumptions. Robustness usually refers to sensitivity to model assumptions. The 95th percentile or even the 99th percentile for distributions that are not normal is not robust at all. One can change the distribution. One can look at a distribution of the same mean and variance and have a 99th percentile very far away from what you would get from the normal distribution.

This is particularly true of insurance losses, if we look at the pure insurance side, if you have a large loss, particularly in property/casualty insurance where very large losses are possible. I would argue very much against the oversimplified model. Now, he did, in all fairness, say this was an oversimplified model; that, in fact, when you look at the probability distributions of numbers and components, they are not lognormal, so they are not normally distributed.

That then begs the question of how to construct the model. One can do it analytically in some cases, but in general, as was mentioned, one builds a model of the various components in effect of the entire company's operations, at least those that are sensitive to short-term changes, and runs with that model. Well, model building is something that actuaries have done for some time, except that they have focused on something different. They have focused essentially on sensitivities, so running individual scenarios will measure sensitivity to those scenarios.

The difference, fundamentally, is that actuaries have not tied probabilities to the different scenarios. The link between the two concepts is that if you tie probabilities to the scenarios, that is, you build a model of interest rate movements, of various other movements, then you essentially do what actuaries have always done in cash-flow-testing models, except with complete probabilistic distributions of all the components.

I have always argued that actuaries should use, in addition to looking at the scenario, stochastic models for all the components in cash-flow testing. Scenario analysis is very important because it is essentially what Mr. Harris mentioned initially—

## DEVELOPING "DOLLARS AT RISK" LIMITS

measuring duration and convexity are essentially sensitivity measures. I think this is the right movement, but what it does for the actuary is really focus on the asset components, as well as the liability components, in building parts of a risk theoretic model, which actuaries have already done.

MR. DOUGLAS S. VAN DAM: I have a couple questions. In the definition of value at risk, it says, "not expected to lose more than that amount of money 95 days out of 100." Would that mean that the amount that you'd lose on any individual day? It's not necessarily the amount that you would lose in the entire 100 days?

MR. HOLTON: Suppose you could relive the next 24 hours 100 times and have 100 different outcomes. Five of those times you would lose more than your value-at-risk number. Then 95 of those times you would lose less than that number, but that is with the portfolio you have right now.

MR. VAN DAM: OK, then the other question is on historical volatility. In terms of re-balancing the portfolio, isn't it possible to look at the historical volatility of the instruments that you've bought and sold and kind of add those together in terms of getting a good estimate of the volatility of the portfolio?

MR. HOLTON: Yes, you could do that and, indeed, what you'd be doing is a value-at-risk calculation.

MR. PETER D. TILLEY: Mr. Harris, when you're dealing with situations and you have a closed period of time, such as a year, and you're using any past gains or losses to calculate the amount that you're going to allow to be put at risk, do you get some really strange-looking discontinuities at the end of the time period? Say you're working on a calendar-year basis and you're halfway through December. You've had a great year. There are all sorts of gains. Do you find that your portfolio managers are dealing with a tremendous amount of funds that can be put at risk for two weeks and then all of a sudden, January 1, boom, they're back to another start date? There tends to be this discontinuity that can cause problems if they're not watching these sorts of things.

MR. HARRIS: I have one thing to clarify. We haven't put this into practice yet, so we don't have experience on that. That's actually an issue that we've discussed. It shouldn't be a big problem with a discontinuity as far as past gains sitting around, because as you recall, the limit remains static for profits. If there had been losses during the year, it's just a tool for managing the risk position a lot tighter toward the end of the year. As the time winds down and many losses have actually been sustained, the risk positions would start to be closed down and processed. That's what I think.

FROM THE FLOOR: One of the problems that we view, because we started to look at the value-at-risk problem at Global Advanced Technology (GAT), is that Monte Carlo simulations are not giving very good accuracy with respect to the confidence interval. Certainly we think that proposing other measures, especially for interest rate risk, such as key rate duration and then rolling that up to perhaps key rate convexity analysis, which is looking at the partial changes, and then the correlations between changes in interest rate movement, might be a better measure to get at the value at risk

in terms of various holding periods that you're looking at. I just want to know if there were any comments on using measures that were certainly better than just parallel sensitivities to get at these things rather than Monte Carlo.

MR. HOLTON: Monte Carlo is a brute force technique. You can make it arbitrarily accurate, understanding, of course, that your volatility estimates may be unstable and that kind of thing. Given your assumptions, you can make the results arbitrarily accurate. People have talked about using duration/convexity-type models for value at risk in which you figure out the duration and convexity of your portfolio and then shock the underliers by one or two standard deviations to see what the price is. That will typically not work.

If Tom Ho was going to be taking that approach, and I have a tremendous amount of respect for the work he does, I'd encourage him not to go down that path. Other people have tried. You get into some statistical problems. Duration convexity works very nicely here if you have one key variable, but having to figure out the upper bound on the 95% confidence interval when you're having to shock multiple key variables, all of them in a duration/convexity correlated sort of manner, leads to problems.

Another severe problem with duration/convexity is that those are factor sensitivities designed typically for delta or gamma hedging or duration/convexity hedging of a portfolio. They describe the sensitivities of the portfolio to very small movements in the underlier. When you start looking at value at risk, you're concerned about how much money you can lose. You're concerned about those infrequent events that occur one day out of 100, or five days out of 100 when you get dramatic market movements. When you start looking at those movements—and that's what you're concerned with in value at risk—duration/convexity can lead to some dramatic errors.

The third release of the J.P. Morgan technical document shows a dramatic example of how, for a simple call option, a duration/convexity, on that particular example a delta/gamma model, actually produces an error larger than a simple delta model. So you can get into some serious problems with that.

FROM THE FLOOR: Yes. I think the J.P. Morgan methodology, though, in just using this pure delta/gamma method, I agree, is very flawed. I think substituting key rate duration sensitivities rather than just a bucket approach actually can get you a much better measure when you're dealing with many complexities. Certainly with Monte Carlo you can miss very large loss occurrences, and this at least gives you, rather than just a tangential surface, the convex surface applied to the various key rate sensitivities.

MR. HOLTON: I think you're right in that if you're going to do a Monte Carlo, you do not want to make your model simply look at parallel shifts in the yield curve. You want to look at multiple points on the yield curve definitely. You could do it from the standpoint of looking at key rate durations, if you wanted to. I view key rate durations as a very interesting tool as an output, as an end result, if you want to compute the key rate durations for your portfolio and have that be your output. As an intermediate step when calculating the ultimate result of value at risk, you can take



## DEVELOPING “DOLLARS AT RISK” LIMITS

that step or you can have a fully robust model that doesn't actually go through key rate durations, but simply considers bends, tilts, and shifts in the yield curve.

FROM THE FLOOR: Are you aware of any other companies, besides yourselves, obviously, that are starting to use this as a management tool and even perhaps going beyond and using it for capital allocation and ultimately for risk-adjusted return-on-capital applications or management?

MR. HARRIS: Many banks are using the application and have used the application, but we do not know of many insurance companies that are using this application. In our development phase, we actually did a survey of different financial institutions. One of the interesting things about that survey—and I think it included about 15 companies, a mixture of banks and insurance companies—was that all respondents said that they would have some form of dollars-at-risk management system implemented within the next two years, and people are thinking about it and starting. I don't think a lot have progressed very far yet in the insurance company environment.

