THE STRATEGIC USES OF VALUE AT RISK: LONG-TERM CAPITAL MANAGEMENT FOR PROPERTY/CASUALTY INSURERS

William H. Panning*

ABSTRACT

In contrast to alternative measures of risk, value at risk (VaR) has important virtues—intelligibility, comparability, and practicality—that make it a potentially valuable tool for strategic decision making and capital management in a wide variety of industries. However, capital-management decisions in most industries—including financial services, such as property/casualty insurance—have time horizons far longer than the one-day horizon that prevails in commercial and investment banking, where the use of VaR is now concentrated. For VaR to be usefully applied to long-horizon decisions, it must address three fundamental problems unique to that context: estimation risk, adaptive risk modification, and franchise risk. This paper describes each of these problems, shows how they can be solved, and provides examples applicable to property/casualty insurance.

The purpose of computing is insight, not numbers.

—Richard Hamming

1. INTRODUCTION

Of the twin concepts, risk and return, that comprise the foundation of modern financial analysis, risk has long been the more troublesome. Return is straightforward both to understand and, in most cases, to measure, despite the persistent efforts of accountants. Risk, by contrast, is elusive, and measures of it, such as standard deviation or beta, fail the critical test of intelligibility: can they be explained in simple terms to one's family or, for that matter, the board of directors? The measures also fail to answer a simple but crucial question. If one alternative has a higher standard deviation or a higher beta than another, just how much greater should its expected return be to justify choosing it? The answer typically offered, which is based upon the aggregate choices of others in the marketplace, exhibits a Panglossian confidence in collective wisdom and the virtue of conformity. Despite these disquieting inadequacies, the textbook measures of risk retain their sway over the generations of MBAs that now populate the world of business.

Value at risk (VaR) offers an alternative way of thinking about, measuring, and managing risk that has three important virtues lacked by other, more traditional measures of risk. First, VaR is intelligible. Roughly speaking, VaR for a firm, a project, or a security is simply the amount of money that it could lose under extremely adverse circumstances. More precisely, VaR is an estimate of the maximum loss that could occur under all but a specified percentage of possible scenarios, ordered from best to worst. Second, VaR is comparable. This potential loss, in dollars or percentages, can be compared across alternatives and bears directly on a classic rule for decision making: never risk more than you can afford to lose. That is, VaR specifies the amount of risk capital a firm needs to undertake a project without threatening its survival. Third, VaR enables us to answer important questions. If one alternative has a higher VaR than another—and therefore requires more risk capital to undertake it—then this additional amount, multiplied by the firm's marginal cost of capital, indicates the

*William H. Panning, Ph.D., is Executive Vice President and Chief Investment Officer at ARM Financial Group, 515 West Market Street, Louisville, Kentucky 40202.

additional return it should require to justify selecting that alternative.[1]

These important virtues make VaR a potentially important tool for strategic decision making in a wide range of situations. During the past decade, for example, VaR has become an indispensable tool in commercial and investment banking, where it is used to determine the amount of capital needed to support a firm's overall risk exposures. If the amount of capital needed exceeds the amount available, the firm reduces its risk exposures by hedging its current portfolio of assets and liabilities. The use of VaR for this purpose is now widespread and has been endorsed by bank regulators. Moreover, banks also use VaR to attribute capital to the component parts of their business, so that they can determine the relative contributions of each to their overall return on capital. This internal application of VaR is therefore an essential tool in banks' continual quest to achieve the maximum return from their capital.

Nearly every firm faces a similar need to estimate its overall need for risk capital and to attribute capital to its component operations or to potential new ones. Consequently, the evident usefulness of VaR in banking strongly suggests its potential applicability in other industries as well. However, as with other important tools, one must calculate and interpret VaR in a manner that is appropriate to the context at hand. It is therefore important to consider whether there are important differences between banks and other firms that need to be taken into account in adapting VaR for more general application.

What distinguishes investment banks from most other institutions is not the particular type of risks to which they are exposed, but rather the time horizon over which information can be gathered and practical decisions implemented. In large investment banks VaR is typically calculated for a one-day horizon, and relevant changes are made to its portfolio of financial assets and liabilities each day to ensure that VaR remains within specified regulatory and internal limits. This short time horizon and rapid response time is made possible by the ability of banks to utilize market information to price their assets and liabilities each day. In most other institutions, by contrast, assets and liabilities can be priced only occasionally—quarterly, yearly, or longer—and risk exposures must be measured and managed over these much longer time spans.

A property/casualty insurer provides an excellent example of the long-horizon nature of some risks and the relevance of VaR to strategic decisions. The balance sheet of a property/casualty insurer consists principally of the items shown in Table 1. On the asset side of the balance sheet are items that can be analyzed by methods currently used in investment banking, since the principal risks are market and credit risk. The liability side, by contrast, is more challenging. The two largest items are the unearned premium reserve and the combined loss and loss adjustment expense reserve (although these are often separated for reporting purposes, I shall here treat them together and refer to them as simply the loss reserve). The economic value (in contrast to the reported value) of the unearned premium reserve is the discounted future loss and expense payments for future accidents covered by in-force policies. By contrast, the loss reserve is an estimate of future payments for past accidents.

For most insurers the loss reserve is the single largest liability on its balance sheet and is the focus of considerable attention by regulators and financial analysts. The existence of this reserve results from the delays between the occurrence of an accident and the submission of a claim, and between the claim report and a final negotiated settlement. For some types of risk both delays can be substantial. Decades may elapse, for example, between an insured's exposure to a toxic substance and the eventual appearance of physical symptoms that trigger a claim. The time lag between provision of coverage and payment of claims can be abbreviated if an insurer switches from occurrence policies, which indemnify the insured against loss from accidents that occur during the term of coverage, to claims-made policies, which pertain to claims actually submitted during the term of coverage. In either case, however, the lag from claim submission to final settlement is often prolonged by legal wrangling.

In determining the value of its loss reserve, a property/casualty insurer must estimate future loss payments both on claims that have been reported but not yet settled and on accidents that have been neither reported nor settled. Because the payments ultimately made may differ considerably from the insurer's current estimate of their value, the insurer must have

[1]Here I have described VaR's application to a choice among mutually exclusive, stand-alone alternatives. By contrast, what matters for a firm with multiple projects or lines of business is the amount of capital needed to support all of them taken together. In this more typical case, what matters in evaluating an alternative is its *marginal* impact on the firm's *overall* VaR. However, VaR has the same virtues in this more complex problem as well.

Table 1
Principal Components of an Insurer's Balance Sheet

Assets	Liabilities
Premiums receivable	Unearned premium reserve
Reinsurance recoverable	Loss and loss adjustment expense reserve
Other receivables	Other payables
Invested assets	Surplus (net worth)

sufficient net worth, known as surplus, to absorb potential adverse deviations from its estimated reserve.

In contrast to investment banks, which can alter their asset and liability exposures rather quickly, the liability risks of a property/casualty insurer change slowly. Property/casualty risk exposures are difficult to trade, so that there is neither an active market nor daily pricing, except in a very few highly specialized circumstances. As a consequence, the principal value of VaR for this industry is strategic rather than tactical. For insurers, it can provide a means to determine their overall need for capital. Moreover, it can assist them in identifying the relative contributions of their various lines of business to their overall capital requirement, so that they can adjust their pricing and business mix to maximize their return on capital. For regulators, VaR provides a potentially powerful tool for ensuring that the companies they oversee are adequately capitalized—a role that it is currently beginning to play in banking as well.

Because of the illiquid nature of insurance liabilities, the strategic decisions just described have time horizons considerably longer than the one-day horizon that characterizes VaR measures in banking. Extending VaR to decisions and alternatives that span these longer time horizons has the potential to enhance its usefulness in nonfinancial industries as well. But extending VaR in this way requires solving at least three important problems that have received little attention. These problems are estimation risk, adaptive risk modification, and franchise risk.

Estimation risk results from the fact that VaR is an estimate of potential loss under an extreme scenario. This estimate is subject to error, since it is based upon a forecast of what asset or liability values will be under conditions that differ from current ones. Estimation error is likely to be quite small for investment banks calculating overnight VaR for portfolios of highly liquid financial assets and liabilities. Applying VaR to other types of assets and liabilities makes estimation error more important. Moreover, extending the time interval over which VaR is calculated magnifies potential changes from current conditions and thus magnifies the impact of errors in forecasting the response of assets and liabilities to such changes. Fortunately, as I shall show, this estimation error can itself be quantified and taken into account in calculating VaR. This is as it should be, since the estimation error is an inherent component of overall risk.

Adaptive risk modification refers to the fact that over an extended period of time the assets and liabilities of a firm do not remain static but are altered in response to implicit or explicit decisions prompted by changing circumstances. In the extreme scenarios that are the focus of VaR, decision makers do not typically remain passive but take actions that alter their perceived risk exposures. Adaptive risk modification is nearly irrelevant for investment banks calculating their overnight VaR, but it can be significant for any firm calculating VaR for longer periods. I shall show how adaptive risk modification can affect VaR and explain how it can be taken into account.

Franchise risk is the potential exposure to loss from assets and liabilities that are not reflected on a firm's current balance sheet. The market value of a firm—the value of its outstanding stock or, for private firms, the amount a potential buyer would pay to acquire it—typically exceeds the economic value of the assets and liabilities included on its balance sheet. What I here call franchise value is this difference between a firm's total market value and its balance sheet value. It consists, in principle, of the discounted expected cash flows from future business. Protecting and enhancing this franchise value is a critical objective at most firms. Consequently, if VaR is to become a longer-term strategic tool, it must take into account potential exposures to loss of franchise value. For example, the discovery of an error in the Intel Pentium chip created only small losses for the firm's balance sheet but had a potentially devastating impact—ultimately averted—on its enormous franchise value from future sales. A VaR calculation that failed to take this effect into account would be incomplete and seriously misleading. The challenge, then, is to measure the components of franchise value and include them in calculations of VaR. In this case I offer only a partial solution, which consists of measuring franchise value for one type of firm and showing how this component of VaR can be calculated for one type of risk.

I begin by presenting a somewhat more thorough introduction to VaR and its differences from other measures of risk. I then devote a separate section to each of the three problems just posed and end with a brief conclusion. Although these problems and their solutions are pertinent to any kind of firm, I present examples applicable to property/casualty insurance. Throughout, I strive to make the exposition widely accessible, although some aspects of VaR are unavoidably technical.

2. VaR VERSUS ALTERNATIVE MEASURES OF RISK

In order to distinguish VaR from other risk measures, it is helpful to consider the specific example shown in Table 2. The table shows four alternatives, each of which has multiple possible payoffs occurring with the different probabilities shown. The four alternatives have the same expected value, but different risks.

Standard deviation, the first risk measure, is a statistical estimate of variability around a mean or expected value, variability that is often portrayed graphically by a bell-shaped curve. For financial assets or liabilities, this measure is calculated from a time series of changes in their value and appropriately adjusted to produce the annualized standard deviation of percentage change in value, the standard measure of return volatility. Despite its widespread use, volatility of return has two important deficiencies as a measure of risk. First, it does not conform to the intuitive definition of risk used by most business professionals. To see why, consider two investment alternatives, one of which (say, one-year Treasury bills or T-bills) will definitely return 5%, while the second has a 50-50 chance of returning 8% or 10%. Although the second alternative has greater variability of return than the T-bill, most investors would not consider it riskier, because its return is always higher. Investors would typically consider the second alternative as riskier only if it had some positive probability of returning less than the T-bill. Their aim is not to achieve certainty, but to avoid loss. Most business professionals thus consider risk as involving a probability of loss—either absolutely (as in a negative return) or relative to some riskless alternative (for example, relative to a T-bill). A second difficulty with standard deviation as a measure of risk is that it can be highly misleading when applied to return distributions that differ dramatically from a bell-shaped curve. In fact, many of the new financial instruments developed in the past decade were specifically designed to create such ''unusual'' return distributions, and their creation has led to the search for improved risk measures.

If losses are what matter, probability of loss could be used as a measure of risk. In Table 2, alternative A has the highest standard deviation of the four alternatives, but the lowest probability of loss. What this measure ignores, unfortunately, is the magnitude of potential losses. It therefore fails to distinguish between two alternatives that have different losses, for example, one small and the other large, but the same probability of loss. Note that in Table 2 alternatives B and C have nearly identical probabilities of loss, 50% and 49%, but that the magnitudes of those losses differ considerably: 20 in one case, 50 in the other.

Expected loss compensates for this deficiency by taking into account both the probabilities and the magnitudes of potential losses. Its value for a given alternative is obtained by multiplying each potential loss by its probability and summing these products. Potential gains are totally ignored. This measure therefore distinguishes sharply between alternatives B

Table 2

Payoffs and Risk Measures for Four Alternatives

	Possible Payoff at Probability				Risk Measures			
Alternatives	50%	49%	1%	Expected Value	Standard Deviation	Probability of Loss	Expected Loss	Worst-Case Loss
A	75	75	2,575	100	249	0%	0	0
B	−20	220	220	100	120	50	10	20
C	249	−50	0	100	149	49	24.5	50
D	104	100	−100	100	20	1	1	100

and C in Table 2 and considers the latter to be far more risky (in fact, the riskiest of the four). Expected loss is the actuarially fair cost of insuring against loss from choosing that alternative, that is, the least amount one could possibly pay someone to compensate for a loss should it occur. In investment terms, expected loss is the minimum cost of a put option with an exercise price of zero. Where the payoff consists of a continuous range of possible values with corresponding probabilities, expected loss can be calculated by an appropriate application of the Black-Scholes (1973) option-pricing formula. For property/casualty insurers, Doherty and Garven (1986), Cummins (1988), and Butsic (1994) have shown how expected loss to policyholders can be calculated and used to determine insurance pricing, fair premiums for state insurance guarantee funds, and minimum capital requirements for insurers.

A final risk measure, one often requested by senior executives, is the *worst-case scenario,* defined as the maximum possible loss. Unlike probability of loss, this criterion finds alternative C in Table 2 to be more risky than alternative B and considers alternative D to have the highest risk of all. In another respect, however, this measure is not very discriminating, for it will fail to differentiate between two alternatives that have the same worst-case loss but dramatically different probabilities of loss. Despite this defect, worst-case loss is a valuable measure for one very important reason. It measures the amount of capital needed to survive the worst outcome that can occur for the alternative being considered. If prudence dictates that one never risk more than one can afford to lose, then worst-case loss informs the decision maker how much capital he or she must be prepared to lose in choosing a particular alternative. As with expected loss, this measure can be generalized to alternatives with a continuous range of payoffs. In such cases one calculates not a worst-case loss (which may be infinite) but a ''highly improbable'' loss. The procedure requires that one rank order the payoffs from best to worst. Then one selects a (very low) probability p of occurrence, 1%, for example, and determines the payoff at the pth percentile of the ordered distribution. If this payoff is a loss of 500, then one infers that in selecting this alternative a capital of 500 is needed to ensure a 99% probability of survival, or, equivalently, to ensure that losses greater than 500 will occur with a probability of 1% or less. If the distribution of possible outcomes is normal or lognormal, then the same value can be directly calculated from the mean and standard deviation of the distribution. Such ''highly improbable'' losses, where the probability is specified, are also called VaR, or VaR$(1 - p)$, where the term $(1 - p)$ is the probability of survival. That is, VaR is a measure of the nearly worst-case outcome.

Generating the Distribution of Outcomes

Since a firm's capital can absorb risks from any of its component operations, VaR analysis is typically applied to the firm as a whole. However, it can equally be applied to a project, a line of business, or a particular asset or liability, depending on one's purposes. In each case, it is calculated from a distribution of possible future outcomes. Most of the work needed to calculate VaR consists of generating that distribution. Here I outline the principal steps in doing so and focus on calculating VaR for the value of a firm whose assets and liabilities consist of fixed-income securities.

1. *Identify and price the components of the firm's value.* Since the objective of VaR analysis is to estimate potential changes in the value of the firm, the starting point must be an estimate of the firm's current value. In investment banking, this is done by determining the market values of the firm's assets and liabilities, which consist principally of securities that are priced daily. Although many of these securities are not actively traded and therefore do not have prices that can be directly observed, their prices are inferred from those of actively traded securities. A similar combination of observation and inference can be utilized in calculating the current value of other types of firms.
2. *Identify the underlying processes that affect outcomes.* The price of a bond and the present value of a loss reserve are determined by the discounted values of their future cash flows, which are in turn a function of Treasury spot rates (zero-coupon bond yields) at different maturities. To model potential changes in the value of a bond or a loss reserve, one must therefore model changes in these underlying Treasury rates.
3. *Estimate the parameters of these underlying processes.* For the bond, one must estimate the volatilities of the spot rates that determine its price. Volatility is the standard deviation (or variance) of changes in each rate. However, changes in spot rates at different maturities are not independent of one another. Rates at different but nearby maturities tend to change in tandem. To take this into account, one must also estimate the covariances among rate changes for different maturities. Since the resulting variance/covariance matrix can be quite large, this process is often simplified in one

of three ways. One is to use fewer rates and interpolate between them as needed. Another is to use factor analysis to reduce the size and complexity of the variance/covariance matrix. A third is to skip parameter estimation entirely and instead use bootstrapping—sampling changes that have occurred in historical data—in the following step.

4. *Generate the distribution of potential changes in the underlying processes.* For a bond or a loss reserve, this is the multivariate distribution of changes in each of the relevant Treasury rates, using the parameters estimated previously. This step is often implemented by means of a Monte Carlo simulation.
5. *For each point in this underlying distribution, calculate the effect on the firm.* For bonds or loss reserves, one must use this distribution of spot rate changes to calculate the distribution of changes in their value. This is done by discounting the relevant cash flows using the combination of spot rates at each point in the multivariate distribution. Doing this for all points in the underlying distribution produces the distribution of overall changes in value.
6. *Select an appropriate value of* p, *and calculate VaR.* As explained earlier, this is done by (a) ordering the various outcomes from best to worst, and (b) selecting the outcome such that only p% of the outcomes are worse. If p is, say, 5%, then this outcome is referred to as the 95% VaR. Its interpretation is straightforward: only $1 - p$ or 5% of the time is the outcome likely to be worse than this value.

The preceding outline is extremely general. More detailed descriptions of VaR calculations for portfolios of financial assets can be found in Jorion (1997) and J. P. Morgan (1995).

3. ESTIMATION RISK

The Nature and Consequences of Estimation Risk

In calculating VaR, as in any statistical modeling effort, there are numerous sources of error or risk that result from our imperfect understanding of the phenomena being modeled. Two that are especially important are specification risk and estimation risk. *Specification risk* results from the use of an inappropriate model that gives wrong or misleading answers. For example, volatility estimates are often based upon an assumed normal distribution of rate changes. By contrast, the actual distribution exhibits a higher central peak and fatter tails than a normal distribution and is more precisely modeled by a t-distribution or by a mixture of normal distributions with different volatilities. Similarly, changes in security prices are sometimes represented as linear functions of rate changes, although the true relationship is typically nonlinear. As in the two instances just cited, specification error can be deliberate, as a means of speeding up calculations, and the magnitude of its effect can be estimated. More commonly, however, it is due to fundamental ignorance. For example, most VaR models are incomplete: the financial market risks incorporated in such models comprise only a portion—although certainly a major portion—of the overall risks to which investment banks are exposed. Losses due to ‘‘rogue traders’’ are one prominent risk not included in VaR calculations. Consequently, banks that use VaR typically compensate for this by using some multiple of VaR to estimate their capital needs.

Even in the absence of specification error, *estimation risk* arises from the fact that the estimated parameters of our models will differ from the true values of those parameters. This type of error arises from the fact that the data employed to estimate parameters are only a sample—a potentially unrepresentative sample—from a much larger set of possible instances.

The following is a graphical example of estimation risk and its consequences. Suppose that the true relationship between two variables, Y and X, is $Y = 2X$, and that we have 100 observations of each, but that our measures of Y are subject to random error. What we observe will not be the true relationship, but rather a scatter of points like those shown in Figure 1.

If we now use linear regression to estimate the relationship $Y = a + bX + e$, where e represents the random measurement error, we will typically not obtain the true values $a = 0$ and $b = 2$, but values that depend upon our particular sample. The greater the random measurement error in Y, the greater the potential deviation of our estimates of a and b from their true values.

Figure 2 shows the various relationships estimated from 12 different simulated samples of 100 observations. Although the range of X in the data was from -5 to $+5$, the range of the X axis in Figure 2 has been increased to better display the differences in the estimates obtained. For 100 different estimates from different samples, the average estimated values of a and b were nearly identical to their true values, but

Figure 1
Illustrative Sampled Data

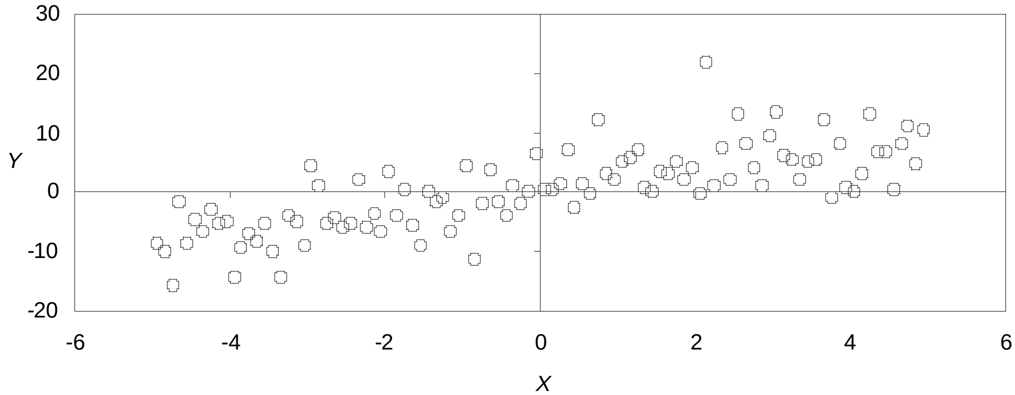

the estimates obtained from any single sample could vary rather widely, despite an R-squared of close to 0.6 in nearly every case.

Especially important in Figure 2 is the fact that the different estimated relationships agree closely at the mean sample values of Y and X and diverge away from that point. As one moves farther and farther away from the average value of the X's we have observed, the greater is the potential for our estimate of the corresponding value of Y to deviate considerably from its true value. Estimation risk is simply this risk that the true value of Y will differ from our estimate of it based on inference or extrapolation from sampled data.

Estimation risk can affect VaR calculations in at least two ways. One is in estimating the effect of changes in market conditions on the values of a firm's assets or liabilities. Among invested assets, mortgage-backed securities provide an especially dramatic example. The cash flows from such securities depend in part upon the proclivity of homeowners to pay off their mortgages prior to maturity, especially when interest rates fall below the rate they are currently paying. Consequently, most algorithms for estimating the prices of such securities rely upon a statistical model for forecasting prepayments for different types of mortgages under different rate scenarios. Such forecasts are subject to estimation risk, which can become particularly severe under the extreme scenarios that are the focus of VaR calculations. However, although estimation risk can be quantified and should be incorporated into these pricing models, it typically is not. Instead, the forecasts are treated as deterministic. This omission can have a serious impact on the estimated price and risk of collateralized mortgage

Figure 2
Relationships Estimated from Different Samples

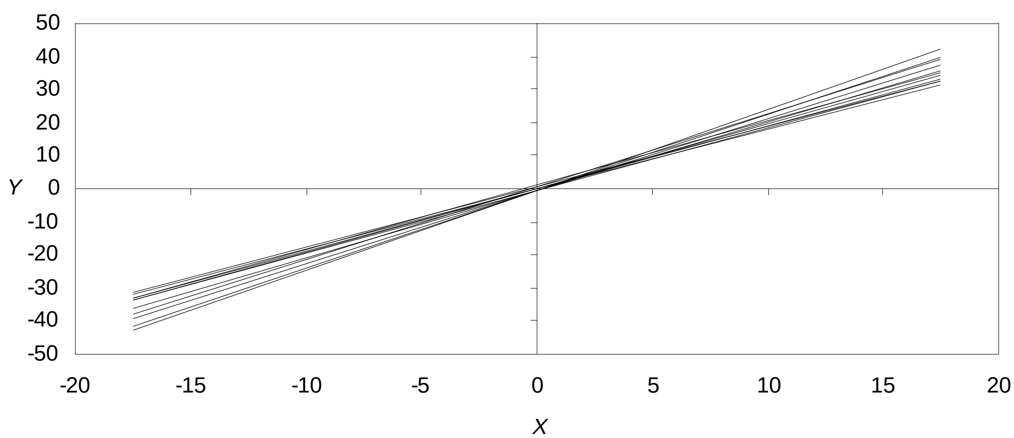

obligations (CMOs), which are securities that give the owner the right to certain classes of cash flows from a pool of mortgages. In particular, CMOs whose cash flows are distant in time have values that are extremely sensitive to estimation error in the prepayment models that are employed. It appears likely, then, that the risks of such CMOs are drastically underestimated because of the current practice of ignoring estimation risk. For some corroborating evidence see Sparks and Sung (1995).

A second important way in which estimation risk can affect VaR is in the estimation of parameters such as the volatility of interest rates or of stock market indices, or the correlations among different types of risk. In estimating these parameters, VaR models typically use daily data from the past. But the volatility estimates that are obtained depend upon the length of the data series employed. Using data from the past several years will produce different volatility estimates—and, consequently, different VaRs—than a longer or shorter data series will produce.

The impact of estimation error or risk on investment decisions is clearly demonstrated by Bawa, Brown, and Klein (1979), who use historical data concerning the risks and returns of various securities to identify efficient portfolios. They show that, in contrast to results obtained when estimation risk is ignored, incorporating estimation risk results in portfolios that are more conservative—that is, less skewed toward assets that have unusually high returns but relatively short track records.

Estimation risk is especially pertinent for VaR calculations that involve non-market risks, that is, risks unrelated to security prices. In such cases the data employed are often limited by availability rather than by choice. Moreover, there may be a substantial lag between the occurrence of events that trigger changes in value and the time at which those changes in value become fully known. During this interval a firm exposed to such risks faces the problem of using the sample of data available from the past to estimate future changes in its value. Estimation risk results from the fact that this sample may be unrepresentative and therefore misleading.

The potential impact of estimation risk on VaR is magnified by the time horizon over which VaR is calculated. When the horizon is long, the potential range of changes from current market conditions is much greater. This, in turn, increases the potential impact of errors in estimating how the prices of securities will be affected by such large changes. This can be clearly seen in Figure 2, which shows that the divergence between the estimated and the true relationship between Y and X increases as one moves away from their mean values in the data used to estimate the relationship. This form of estimation risk is further compounded by errors in estimating market volatility. The longer the time horizon over which VaR is calculated, the greater the potential impact of this type of error as well.

Incorporating Estimation Risk in VaR

A property/casualty insurer's loss reserve poses two important statistical problems. The first is estimating the expected value of future loss payments. The second is estimating the distribution of possible deviations from this expected value and the consequent amount of surplus the firm needs to absorb adverse deviations with a high probability of remaining solvent. The first of these problems has been studied intensively by actuaries. The second, by contrast, has received far less attention. In practice, surplus requirements are still widely determined by two traditional rules of thumb. One focuses on the ratio of insurance premiums to surplus, the other on the ratio of reserves to surplus. Both are inadequate to the task. The first really focuses on the risks posed by writing new business, rather than the risks posed by loss reserves. The second is clearly pertinent but is typically based upon historical industry averages rather than upon a rigorous quantitative foundation. Both are incomplete measures, since the surplus needed by an insurer depends upon its total risk exposures, of which adverse reserve development is only one component. Here I shall focus on the twin problems of estimating the expected value of an insurer's loss reserve and determining the distribution of potential deviations from this expected value. Solutions to these two problems are essential if VaR is to become a useful strategic tool for property/casualty insurers. Moreover, a correct solution to both problems requires attention to the problem of estimation risk.

Two caveats are in order. First, adverse loss reserve development is only one of the numerous risks that affect an insurer's overall VaR and need for surplus. Consequently, the analysis presented here is intended to be illustrative rather than complete. Second, loss payment data often exhibit complexities not present in the data used here. Although the model I present in Table 3 can be elaborated to take such complexities into account, addressing these issues here would detract from the principal point of this section: the fact

Table 3
Loss Development Triangle

Accident Year	Development Year									
	0	1	2	3	4	5	6	7	8	9
0	1,291	791	548	423	334	304	224	230	214	210
1	1,492	881	617	471	390	282	263	241	231	
2	1,640	944	685	546	379	320	295	258		
3	1,818	1,049	799	522	409	364	297			
4	1,992	1,171	774	566	482	376				
5	2,056	1,133	814	644	476					
6	2,004	1,149	886	609						
7	2,023	1,264	824							
8	2,131	1,203								
9	1,998									

Source: Adapted from Corporate 10K.

that long-term VaR must take estimation risk into account, and in a way that is systematic rather than ad hoc.

The first step in calculating VaR for a loss reserve is to determine the expected future cash flows for past accidents. The data used by property/casualty insurers and regulators to do this are shown in Table 3, here for a large firm that writes long-tail business, that is, lines of business in which loss payments extend over a long period of time. The data shown here have been adapted from the firm's 10K in two ways. First, all numerical values have been multiplied by a constant, to conceal the firm's identity. Second, although the original data were cumulative, from left to right, I have here shown the decumulated values. The use of decumulated values is essential for two reasons. First, they are the actual payments that occurred in each period and therefore constitute the correct basis for estimated expected cash flows in each future period. Second, I will model these payments as consisting of two parts: an expected or predictable part, and a random deviation or error. It is these deviations or errors that will enable us to ultimately calculate the distribution of overall deviations from the expected value of the reserve. The use of cumulative values, which is common, conflates errors for different periods and makes them difficult to model correctly.

Each row in Table 3 represents a particular *accident year,* defined as the year in which an accident occurred that was covered by the firm's policies. Here I have replaced the actual calendar year numbers with the digits from zero to nine. Each column in the table represents a particular *development year,* defined as the number of years subsequent to the accident year in which payments were made. A payment occurring in development year zero was in fact made in the same calendar year as the one in which the accident occurred; payments in development year one occurred in the calendar year following the accident; and so on. The entry P_{ij} in row i and column j of the table is the total amount paid by the firm for losses incurred in accident year i but paid in development year j, that is, j years later. Note that the payments on any diagonal of the table were made in the same calendar year $i + j$.

The firm's loss reserve is an estimate of the total payments it will make in future years for accidents that have already occurred. These payments consist of the unknown entries in the blank portion of the table, as well as those occurring to the right of the table, in development years ten and higher.

Actuaries have developed a variety of methods for estimating loss reserves. Unfortunately, the methods most commonly in use are ill-suited for determining VaR. There are two important reasons for this. First, these methods are essentially ad hoc, in that they lack a basis in an underlying model of the data. As a consequence, they produce parameters without reference to a well-defined statistical measure of goodness of fit to the data. Second, because of this they provide no statistical basis for estimating the magnitudes and probabilities of possible deviations from the estimates they produce. The procedure employed here, which remedies both these defects, is based upon work by Taylor (1987).

The model used here to represent the data in Table 3 is

$$P_{ij} = a \cdot f(i) \cdot f(j) \cdot e_{ij},$$

where a is an estimate of P_{00}, $f(i)$ is a function that represents changes in the volume of losses over accident years, $f(j)$ is a function that represents the pattern of payouts over development years, and e is an error term. For example, suppose that losses are growing by a constant percent each accident year. In that case $f(i) = \exp(bi)$, where b is the growth rate and i is the accident year. Similarly, if loss payments in each year are decreasing at a constant rate, then $f(j) = \exp(cj)$, where j is the development year and c is the percentage change in payments from one development year to the successive one. More complex functions can also be used if needed. For example, we might use $f(j) = \exp(c_1 \cdot j + c_2 \cdot j^2 + c_3 \cdot j^3)$.

By taking logarithms of both sides of the preceding equation we obtain

$$\ln P_{ij} = \ln a + \ln f(i) + \ln f(j) + \ln eij,$$

which is in a form suitable for use in linear regression. For the data in Table 3 we will here specify $f(i)$ as $\exp(b_1 \cdot i + b_2 \cdot i^2)$ and $f(j)$ as $\exp(b_3 \cdot j + b_4 \cdot j^2 + b_5 \cdot j^3)$. The final form of the equation that we will use is therefore

$$\ln P_{ij} = \ln a + b_1 \cdot i + b_2 \cdot i^2 + b_3 \cdot j + b_4 \cdot j^2 + b_5 \cdot j^3 + \ln e_{ij}.$$

The parameter estimates obtained from linear regression are shown in Table 4. In addition to the estimated values of the model parameters, Table 4 also shows the standard errors of each parameter (the standard deviation of the estimated parameter value) and the corresponding t-statistics, the ratios of the estimated values to their standard errors. A t-statistic with an absolute value greater than 2.0 indicates that the estimate obtained is highly unlikely to have a true value of zero. Overall fit to the data appears to be excellent, as indicated by a high value of R-squared and a low standard error for the overall equation.

This apparent excellent fit between model and data is confirmed by several additional comparisons between the data and the estimated values. Table 5 shows the estimated values of P_{ij} obtained from the regression equation.

The error terms in the model, consisting of differences between the logged original data and the fitted logged values, are shown in Table 6. These are approximately equal to the percentage differences between fitted and actual unlogged values. One would statistically expect one or two of these values to be at least twice the standard error of the equation, and that is indeed the case. However, there appear to be no systematic patterns that would suggest possible specification error.

The absence of specification error can be confirmed in another way as well. Systematic accident year errors can be detected by calculating and comparing, for the

Table 4

Parameter Estimates for Loss Payment Data with ***R***-Squared 0.996 and Standard Error 0.047

Parameter	Estimated Value	Standard Error for Parameter	t-Statistic
$\ln a$	7.199	0.020	360.9
b_1	−0.526	0.017	−30.8
b_2	0.049	0.005	9.5
b_3	−0.0016	0.0004	−3.7
b_4	0.114	0.009	13.2
b_5	−0.008	0.001	−7.1

Table 5

Fitted Loss Payment Data

Accident Year	Development Year									
	0	1	2	3	4	5	6	7	8	9
0	1,339	830	562	413	325	272	240	221	210	205
1	1,489	923	625	459	361	303	267	246	234	
2	1,630	1,010	685	502	396	331	292	269		
3	1,758	1,089	738	542	427	357	315			
4	1,866	1,156	783	575	453	379				
5	1,950	1,208	819	601	473					
6	2,007	1,244	843	618						
7	2,034	1,261	854							
8	2,030	1,258								
9	1,994									

Table 6
Deviations of Actual from Fitted Loss Payment Data

Accident Year	Development Year									
	0	1	2	3	4	5	6	7	8	9
0	−3.5%	−4.7%	−2.5%	2.5%	2.9%	11.0%	−7.0%	4.1%	2.0%	2.5%
1	0.3	−4.5	−1.2	2.7	7.7	−6.9	−1.4	−1.9	−1.2	
2	0.7	−6.7	0.2	8.4	−4.1	−3.4	1.1	−4.1		
3	3.5	−3.7	8.1	−3.5	−4.2	1.9	−5.7			
4	6.7	1.4	−1.1	−1.4	6.4	−0.8				
5	5.4	−6.4	−0.4	7.1	0.6					
6	0.0	−7.8	5.1	−1.5						
7	−0.4	0.4	−3.5							
8	5.0	−4.4								
9	0.3									

logged actual and fitted values, the difference between the average (logged) values in each row and the average of the corresponding values in the top row. Systematic development year errors can likewise be spotted by applying the same procedure to the columns of the logged actual and fitted data. These comparisons are shown in Figures 3 and 4. Both confirm a close fit to the data and the apparent absence of specification error in the model.

Finally, the distribution of errors, fitted here by means of a procedure called kernel density estimation[2] (Silverman 1986), displays symmetry and a reasonable approximation to a normal distribution, as shown in Figure 5.

The results obtained so far now enable us to obtain the expected future loss payments that comprise this insurer's loss reserve. These forecast values are subject to three sources of error:

1. *Specification error:* systematic differences between the model and the data on which it is based, differences that will produce a biased forecast. The extensive analysis of errors and comparisons between actual and fitted values were designed to ensure that this source of error has been minimized.
2. *Random error:* random deviations between actual and fitted or forecast payments. These errors are represented by the error term in the fitted equation and can be anticipated to occur in the future as well. Their magnitude is estimated by the standard deviation of the fitted model. Note, however, that this model applies to the logarithms of the data. Consequently, actual errors will in fact be distributed lognormally for each payment.
3. *Estimation error:* differences between the true values of the model parameters and the estimated values obtained from the data. This source of error is reflected by the fact that each parameter has a standard error that represents the probability distribution of this actual value around its estimated value.

All three sources of error affect the values of forecast future payments. Because the independent variables—here consisting of powers of i and j—are known for each future payment, we can calculate the logged forecast payments using the parameters estimated from the available data. However, the expected unlogged payment is not the antilog of the expected logged payment. Instead, it is $\exp(E_{ij} + S_{ij}^2/2)$, where E_{ij} is the expected logged value of payment P_{ij} and S_{ij} is the standard deviation of the forecast error of that logged value. Forecast error reflects both random error and estimation error.

The variance of the forecast error for the individual (logged) future payments is the main diagonal of the variance/covariance matrix of forecast errors. This matrix, $\mathbf{V}$, is equal to $se^2[\mathbf{I} + \mathbf{X_f}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X_f}']$, where se is the standard error of the regression equation, $\mathbf{X}$ is the matrix of independent variables used in the regression (which includes a leading column of ones for estimating the intercept term), $\mathbf{X_f}$ is the matrix of independent variables for the unknown logged payments to be forecast, and $\mathbf{I}$ is the identity matrix, which has ones on its main diagonal and zeros elsewhere.

The resulting (unlogged) forecast future loss payments through development year nine are shown in

[2]Kernel density estimation, a procedure for producing a relatively smooth distribution from N data points, consists of the following steps: (1) Create N normal distributions, each with a mean equal to one of the data values, and a standard deviation S. (2) Calculate the overall distribution by averaging the N individual distributions. (3) Use one of several criteria (such as maximum likelihood) to determine an appropriate value of S. When S is tiny, the result will consist of N spikes. As S is increased, the number of spikes will diminish until the overall distribution ultimately becomes nearly flat.

Figure 3

Actual and Estimated Log Accident Year Payment Ratios

Figure 4

Actual and Estimated Log Development Year Payment Ratios

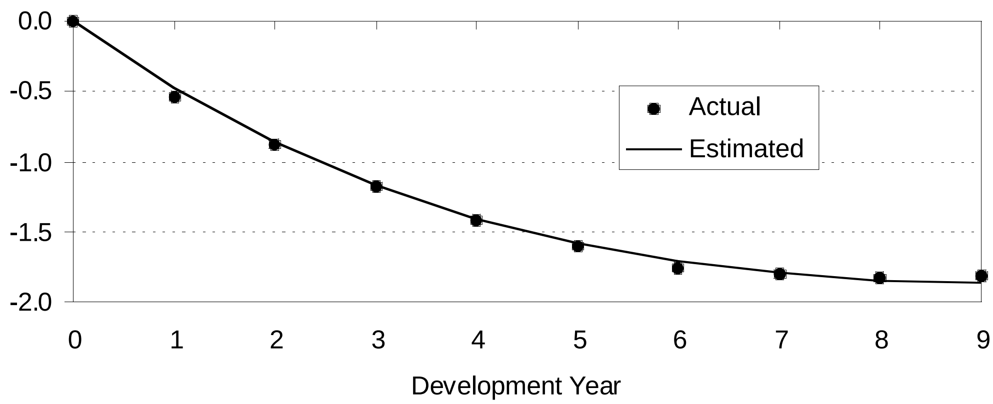

Figure 5

Fitted Error Distribution

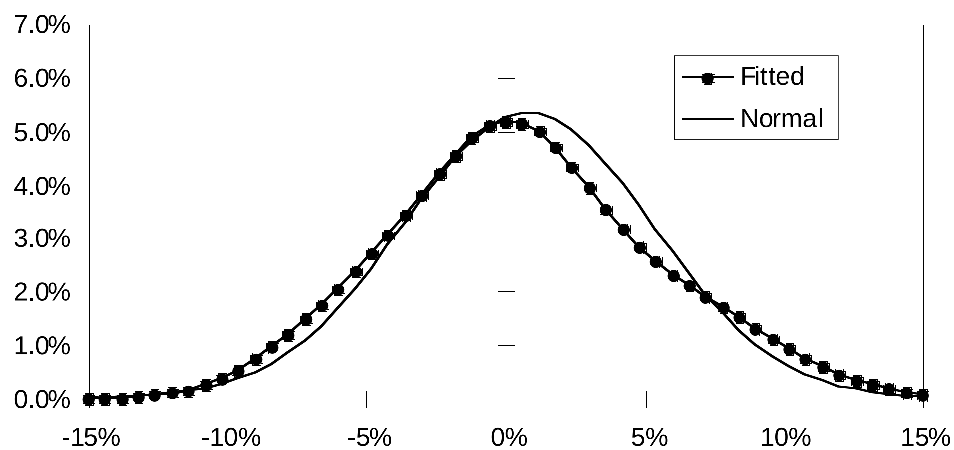

Table 7, and their standard errors (as percentages) are shown in Table 8.

It is evident from the forecast payments in Table 7 that additional payments are likely to occur in development years ten and above. Here these payments were forecast by assuming that the proportional decrease in value from one development year to the next is identical to the decrease from development years eight to nine in the fitted values in Table 5. For example, in the first row of Table 5, the forecast values for development years nine and eight are 205 and 210. The ratio 205/210 = 0.975. Consequently, the forecast values for all development years greater than ten are $E[P_{i9}] \times 0.975^{(J-9)}$.

This procedure makes it possible to develop an overall development year payout pattern, which consists of the percentage of total accident year losses that are paid out in each development year. This payout pattern for the company analyzed here is shown in Figure 6.

We have now obtained estimated loss payments $E[P_{ij}]$ for all future years. From these estimates we can calculate the estimated undiscounted loss reserve, which is $\Sigma E[P_{ij}]$ for all $j > i$. We can also discount each such cash flow at an appropriate rate to obtain a discounted value for the reserve.

The procedure just described solves the first problem: establishing the expected value of the insurer's loss reserve. What remains is the second problem: determining the distribution of possible deviations from this expected value. This is an essential step in calculating VaR for the loss reserve or for the firm as a whole. The key in both cases is **V**, the variance/covariance matrix of forecast errors. Inspection of this matrix shows that forecast errors are in fact correlated with one another. That is, one cannot assume that the forecast errors for different future payments are independent of one another. Fortunately, the interdependencies among forecast errors can be taken into account by means of Monte Carlo simulation. Here I used @RISK™ to generate 1,000 scenarios, in each of

Table 7
Estimated Future Loss Payments

Accident Year	Development Year									
	0	1	2	3	4	5	6	7	8	9
0										
1										228
2									256	250
3								290	276	269
4							334	308	293	286
5						396	349	322	306	298
6					487	408	360	331	315	307
7				627	494	413	365	336	319	311
8			852	625	493	412	364	335	319	311
9		1,236	837	615	484	405	357	329	313	305

Table 8
Standard Deviation of Estimated Future Loss Payments

Accident Year	Development Year									
	0	1	2	3	4	5	6	7	8	9
0										
1										6.0%
2									5.2%	6.1
3								5.0%	5.2	6.2
4							5.0%	5.0	5.3	6.3
5						4.9%	5.0	5.1	5.3	6.3
6					4.9%	5.0	5.0	5.1	5.4	6.3
7				5.0%	5.0	5.1	5.2	5.2	5.4	6.3
8			5.2%	5.2	5.3	5.3	5.4	5.4	5.6	6.4
9		5.5%	5.6	5.7	5.8	5.8	5.9	5.9	6.0	6.7

Figure 6
Development Year Payout Pattern

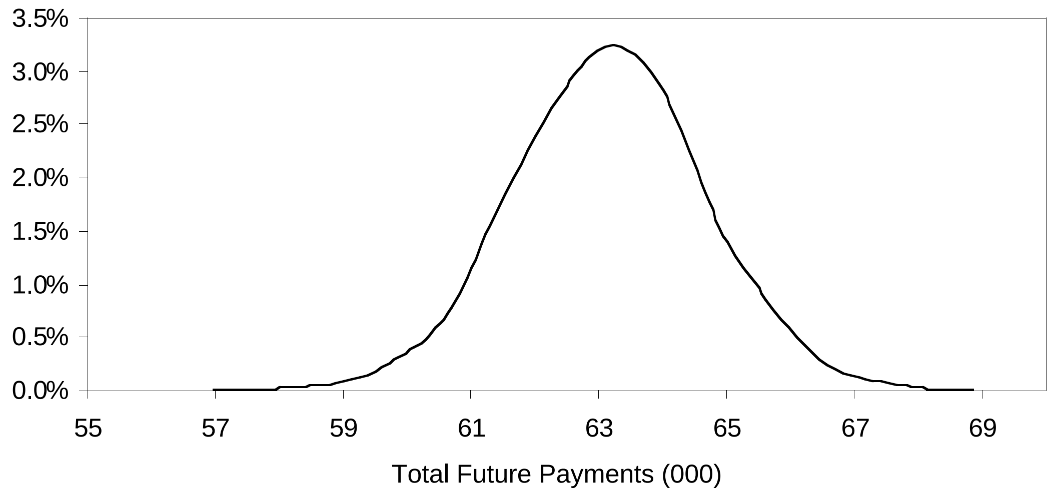

Figure 7
Distribution of Total Nominal Payments

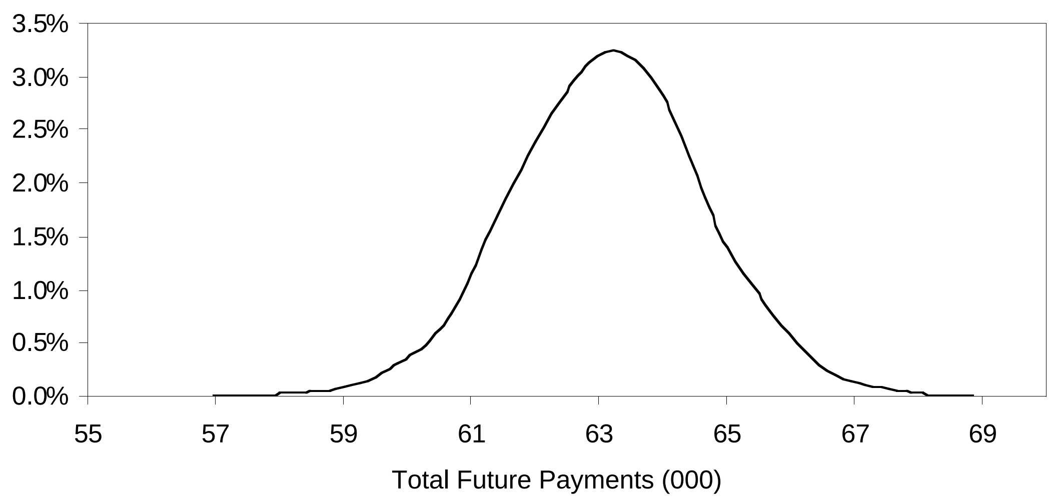

which a random error was selected for each forecast future payment from a multivariate distribution that reflected the variances and covariances of forecast errors. For each scenario I calculated two values: the sum of the future payments, and their discounted value at a 6% annual rate of interest. The distributions of the nominal and discounted totals are shown in Figures 7 and 8.

Summary statistics for the two distributions are shown in Table 9. The first column shows the mean of the two distributions, and the second column shows the first percentile for each distribution. The difference between this second number and the first, shown in the third column, is the 99th percentile VaR in each case. This is the amount of surplus the firm should have in order to be 99% confident of being able to successfully fund its loss reserve liabilities. The final column in Table 9 shows the ratio of the expected payments to this surplus. These ratios are quite high relative to industry averages and widely used rules of thumb. However, a portion of this difference is no doubt attributable to the fact that other sources of risk not examined here require additional surplus.

The objective of this example has been to demonstrate how estimation risk can be incorporated into the calculation of VaR. It is not intended to provide a complete analysis of the risks and capital requirements posed by an insurer's loss reserve. It takes no

Figure 8
Distribution of Total Discounted Payments

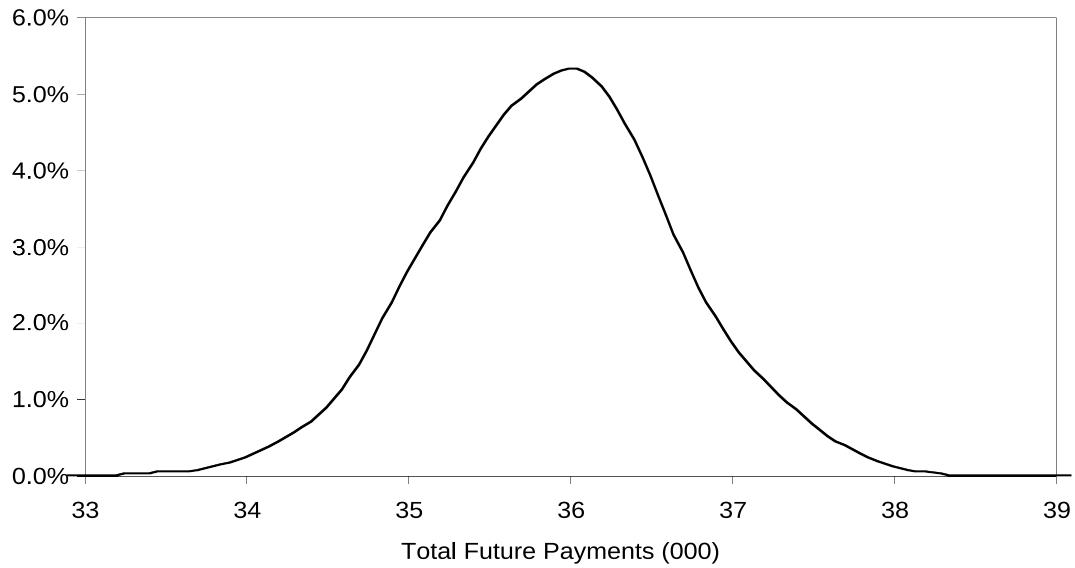

account, for example, of potential changes in the value of the discounted reserve resulting from changes in interest rates. This risk has been extensively discussed elsewhere (see, for example, Panning 1995) and is highly correlated with concomitant changes in the market value of the insurer's invested assets.

The overall significance of estimation risk for VaR estimates depends upon the length of the time horizon, as noted earlier. For the particular case of loss reserves, it also depends upon both the loss payment pattern and the magnitude of random error to which individual payments are subject. Short-tail lines of business, in which most losses are paid within a relatively few years, are subject to less estimation risk than are long-tail lines. This is clearly seen in Figure 9, which shows the magnitude of forecast error by development year. It is clear that forecast error begins to grow exponentially beyond year nine. Since forecast error consists of random error and estimation error, and since the former is roughly constant, the upward trajectory of the graph is due almost solely to increases in the magnitude of estimation risk.

The significance of estimation risk also increases with the magnitude of random error. The random component of forecast error varies across different lines of business and across different insurers. When random error is small, as in this example, estimation risk is also relatively small. But as random error increases (such as across different lines of business), so does the magnitude of estimation risk. Nonetheless, as this example has demonstrated, the magnitude of estimation risk can be quantified and included in the determination of VaR.

4. ADAPTIVE RISK MODIFICATION

In calculating VaR for an extended time horizon, one cannot safely assume that the assets and liabilities of a firm remain static. Even under the most benign scenarios, changes will occur because of normal business operations. Under the extreme scenarios that are the particular focus of VaR calculations, however, dramatic changes will almost certainly occur as decision

Table 9
Summary of Monte Carlo Simulation Results

Simulation	Mean	First Percentile	Difference	Reserve to Surplus Ratio
Nominal	63,120	66,560	3,440	18.3
Discounted	35,863	37,600	1,737	20.6

Figure 9
Forecast Error by Development Year

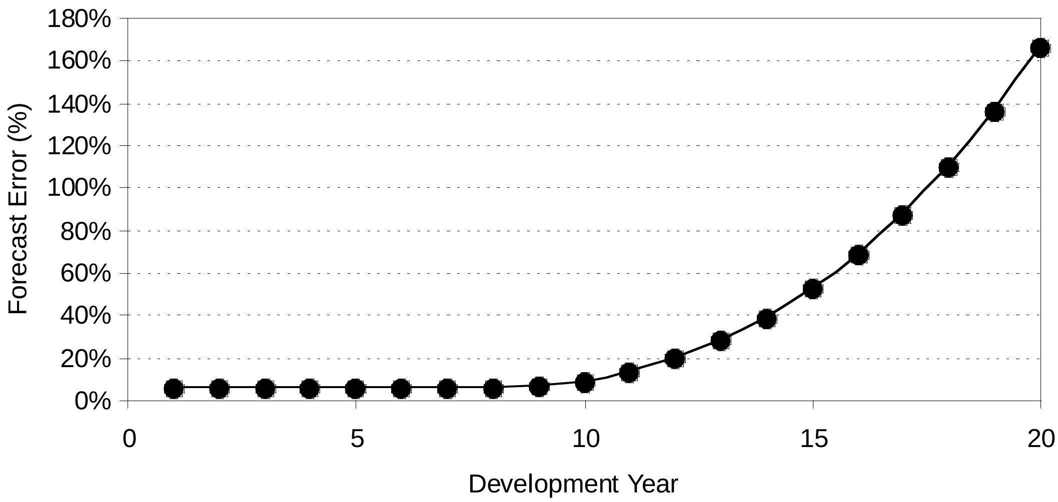

makers attempt to modify their risk exposures in response to the changing circumstances of the firm or of their particular portion of it.

An excellent example of this phenomenon is found in most investment operations. With the exception of managers of index funds, most traders and portfolio managers seek to enhance investment returns by actively trading their accounts in response to changes in the marketplace. In some cases their responses are planned and explicit. In others they are instinctive and implicit. In nearly all cases, however, their actions have systematic consequences for VaR calculations.

To illustrate this point in some depth, I simulate the behavior of three investment portfolios that are initially identical; that is, they have the same initial market value (100) and the same initial asset composition, with 50% invested in one-year bonds yielding 5% and 50% invested in a non-dividend-paying stock that has an initial price of 100, an expected return of 10%, and a standard deviation of 20%. A VaR calculation that assumed these portfolios would remain static would produce the same number for all three portfolios. In fact, one of the three portfolio managers follows a buy-and-hold strategy and does indeed have a static portfolio composition. However, the other two portfolios do not remain static over time. One portfolio manager is a momentum buyer, who buys additional stock as it rises in price and sells stock when its price declines. A second portfolio manager follows the classic ''buy-low-sell-high'' strategy and consequently buys stock when its price falls and sells stock when its price rises. In the simulation, their response to changing stock prices is equal but opposite in direction, so that their combined portfolios have the same static composition as the buy-and-hold portfolio.

Each scenario in this simulation consists of 250 market days of activity, which is equivalent to a year of calendar time. Each day the bonds held by each portfolio accrue interest, and the stock changes in price in a logarithmic random walk: that is, the daily percentage changes in the price of the stock comprise a normal distribution in which price changes on successive days are independent. At the end of each day, the two active managers buy or sell stocks and bonds in response to these price changes. Transaction costs are assumed to be zero.

The rules followed by the two active managers are identical in form:

$$P = \frac{1}{1 + \exp(a \cdot b - b \cdot c)},$$

where P is the percentage of the portfolio invested in stock, a and b are parameters, and c is the cumulative percentage change in the price of the stock since inception (the beginning of the simulated year). The parameter $a = 0$ for both managers. The parameter b determines the direction and magnitude of their response to changes in the stock price. I have set $b = 10$ for the momentum manager, and $b = -10$ for the buy-low-sell-high manager. A graphical illustration of the response curve for the momentum manager is shown in Figure 10. The response curve for the other manager is the left-to-right mirror image of the one

Figure 10
Momentum Response Curve

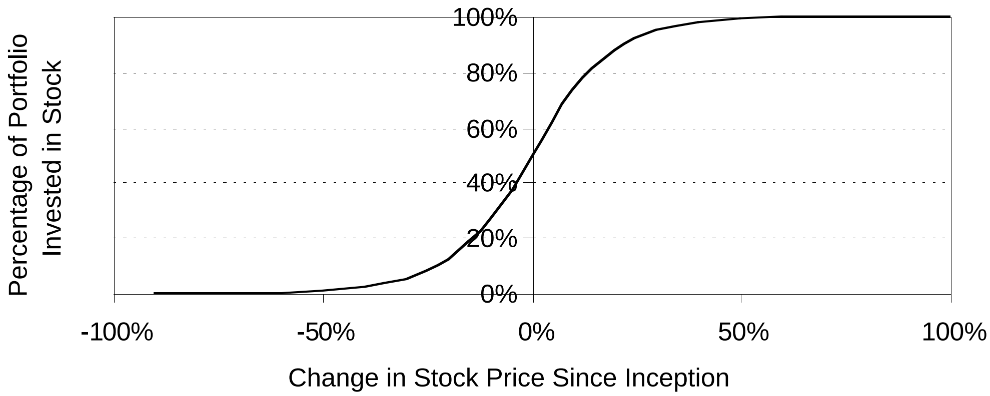

shown. At the midpoint of the response curve, the response to stock price changes is approximately $b/4$. That is, for the momentum manager, a 1% change in the stock price will induce him or her to increase the stock percentage of his portfolio by 2.5%, relative to the proportion held the previous day. Note that a portion of this increase will have occurred naturally due to the increase in the price of the stock.

For each yearlong scenario, consisting of 250 daily price changes, the ending values for the three portfolio managers will typically be different. For 1,000 such scenarios, the distribution of year-end portfolio values for the three portfolios is shown in Figure 11.

The distribution for the buy-and-hold portfolio is lognormal. Relative to it, the momentum strategy clearly has a longer upside tail and a lower downside risk, and the buy-low-sell-high strategy has just the opposite characteristics.

Further insight into the consequences of the two active strategies can be gained from Figures 12 and 13, which show the relationship between the ending value of each portfolio and the ending price of the stock for each of the 1,000 scenarios. The straight line in both graphs shows the ending value of a portfolio consisting only of stock. The results for the momentum strategy closely resemble those obtained by owning the stock and buying a put option on the stock with an exercise price of about 90. By contrast, the results for the buy-low-sell-high strategy resemble the results obtained from owning the stock and selling a call option on the stock with an exercise price of about 110.

As one would expect from these graphs, the momentum strategy has the lowest VaR of the three portfolios, and the buy-low-sell-high strategy has the highest. The results for one set of 1,000 scenarios are

Figure 11
Return Distributions for Alternative Strategies

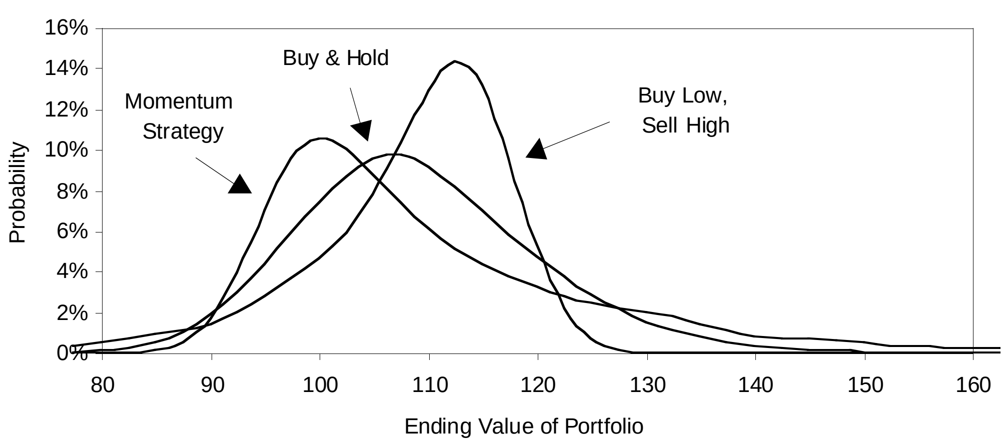

Figure 12
Results of Momentum Strategy

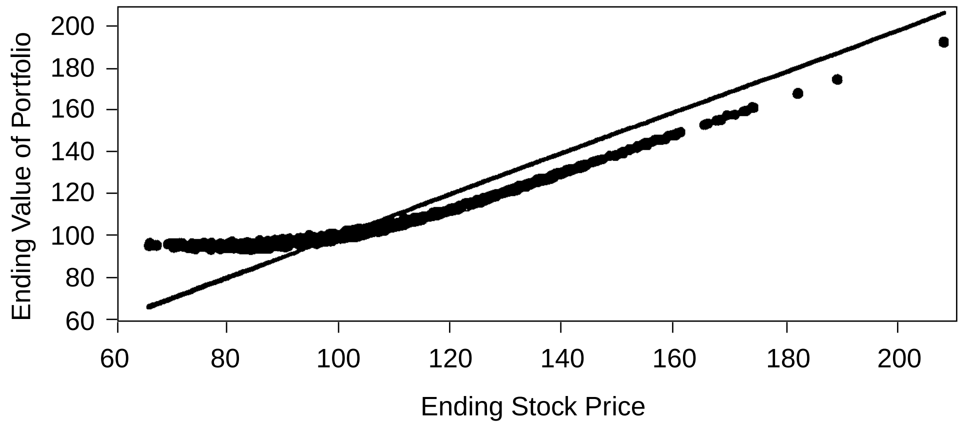

summarized in Table 10, in which VaR is expressed as a percentage of the initial portfolio value.

These results show that, when VaR is calculated for a long-term horizon, one cannot simply examine the initial composition of a portfolio of assets or liabilities. Instead, one must include in the VaR calculation the likely response of asset or liability managers to the events included in the scenarios that serve as the basis for the VaR calculation. While this may seem difficult or even impossible, in fact it presents a valuable opportunity, since different strategies have quite different implications for the capital needed by the firm. The real challenge is to make the strategies that decision makers follow explicit rather than implicit, to systematically examine their implications for the firm, and then to ensure that these strategies are in fact followed. Doing this can have substantial benefits for the firm. For example, it is unlikely that most investment officers realize that following a buy-low-sell-high strategy requires a considerably higher commitment of the firm's capital than following a momentum strategy, quite apart from differences in the average returns of the two strategies.

This conclusion applies not only to investment portfolios, but to any portfolio of assets and/or liabilities. For example, many firms follow an implicit momentum strategy with the various lines of business that they manage, a strategy that consists of investing to expand a business when its profitability is high, and contracting or selling it when its profitability declines. Using such a strategy creates what are called ''real options,'' asymmetric return profiles similar to the

Figure 13
Results of Buy-Low-Sell-High Strategy

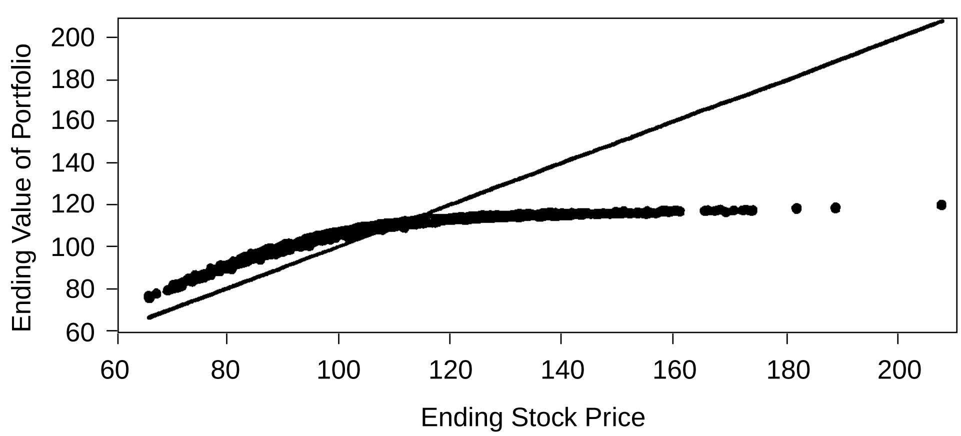

Table 10
VaR for Three Portfolio Strategies

Percentile	Momentum	Buy and Hold	Buy Low Sell High
95%	5.3%	6.6%	8.4%
99	6.2	11.8	18.7
99.9	7.1	14.5	24.7

one produced by the momentum strategy just presented. This subject is treated in depth by Dixit and Pindyck (1994) and Trigeorgis (1996).

5. FRANCHISE RISK

The objective of most strategic decisions is to increase the value of the firm. VaR is potentially useful in that context because it focuses on the marginal implications for risk capital of the different alternatives under consideration. But if VaR is to become a useful strategic tool, it must be able to take into account the various risks to which the value of the firm is exposed. This, in turn, implies that VaR must be based upon a correct initial estimate of the firm's total value.

In the case of a property/casualty insurer, an obvious procedure for estimating the value of the firm would be to ascertain the market value of the firm's assets and subtract from that the estimated present value of the firm's liabilities, suitably adjusted for their risk. This is in fact a common practice in estimating and managing the sensitivity of an insurer's surplus to changes in interest rates. Unfortunately, this procedure is highly misleading, for it ignores that portion of the firm's value that cannot be ascertained from its balance sheet.

The market value or purchase price of an insurer is typically higher than the amount obtained by economically valuing its balance sheet. The difference consists of what I here call *franchise value:* the present value of the firm's expected cash flows from business it will write in the future. Most of an insurer's clients tend to renew their policies, since there are significant practical costs to switching to another firm. Moreover, insurers create incentives for their clients to renew. As a consequence, a typical insurer can expect to retain around 60–90% of its existing clients from one year to the next. Even if it attracts no new clients, an insurer can anticipate significant future profits from retentions alone. Unfortunately, accounting standards do not permit this franchise value to be recognized unless a firm is in fact purchased, in which case the excess of the purchase price over book value is treated as an asset and inserted onto the balance sheet of the acquirer.

This fact has two important implications for the strategic use of VaR. First, franchise value is too important to be ignored, since enhancing that value is an important objective of many strategic decisions. Second, the components of franchise value, and their exposures to various types of risk, should be modeled and included in VaR in the same fashion as components that are in fact recognized by accounting standards. In this section I show how this can be done for a property/casualty insurer with respect to one source of risk. A more detailed treatment of this problem is presented in Panning (1995).

In the example presented here I focus exclusively on franchise value and the potential changes in its value due to changes in interest rates. I have simulated an insurer that currently writes policies with a premium of 100 and loss payments totaling 100 with a payout pattern that decreases gradually from 25% to 2% in development years zero through ten. Its retention rate is 90%, and the current interest rate is 6%.[3] All other complexities, such as expenses, taxes, and changes in loss ratios with subsequent renewals, are ignored for purposes of illustration. The present value of its profits from the firm's retentions (I ignore potential sales to new clients) is approximately 83.6.

If interest rates change but all other conditions including premiums remain constant, the present value of these retentions—the firm's franchise value—will change accordingly. Figure 14 shows the effect of different rate changes on the franchise value. Note that franchise value increases when rates rise: behavior that is opposite from that of a bond. This is due to the fact that the duration (interest rate sensitivity) of future premiums is lower than the duration of future loss payments (which occur later in time than premium inflows).

[3]This is not the firm's cost of capital. The fact that a firm must provide a reasonable rate of return to its suppliers of capital has strong implications for the profit margins it must generate, but virtually no bearing on the appropriate interest rate for discounting its future cash flows.

Figure 14
Effect of Interest Rate Changes on Franchise Value
When Premiums Remain Constant

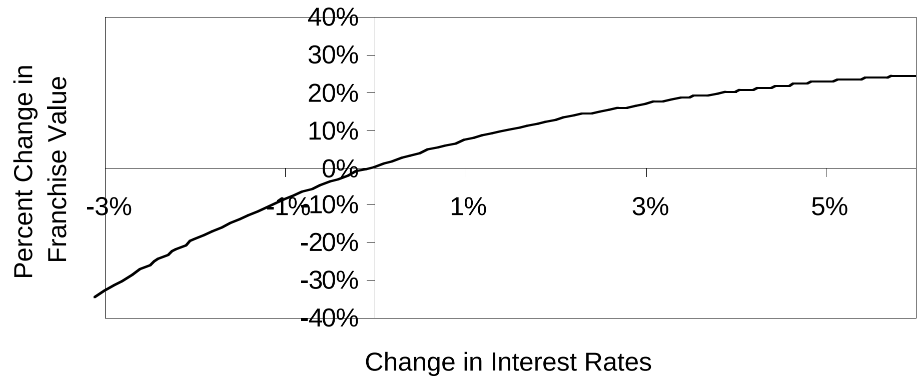

Since interest rates do in fact change, I have simulated the distribution of changes in franchise value that results when the volatility (standard deviation of interest rate changes, as a percentage of the initial rate) is 15%. This distribution, shown in Figure 15, is reverse lognormal, as one would expect from Figure 14. The 99% VaR for franchise value is 21% of its initial value.

Now let us add one small but important—and realistic—complexity to the simulation. Let us assume that, when interest rates rise or fall, at least some firms in the insurance marketplace alter their prices (premiums) to gain or retain market share. If the firm we are simulating maintains its current price, then its retention rate will decline, since at least some of its clients will be attracted to the competitors' lower prices. On the other hand, if it matches its competitors' price cuts, then its franchise value will be lower than would otherwise be the case.

In either case the overall effect is nearly the same: lower cash flows from retentions. The results for this more complex case are shown in Figure 16. Note that the effect of interest rate changes on franchise value is now reversed: an increase in interest rates now reduces the firm's franchise value.

The distribution of possible changes in franchise value for this situation is shown in Figure 17. The 99% VaR is again 21% of the initial franchise value, but this now occurs under conditions that are the opposite of those found in the preceding simpler case.

Figure 15
Distribution of Changes in Franchise Value
When Premiums Remain Constant

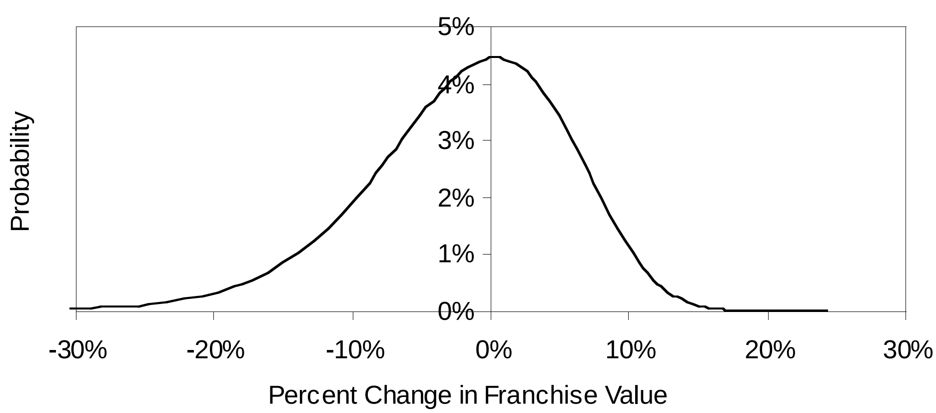

Figure 16
Effect of Interest Rate Changes on Franchise Value
When Premiums Vary Inversely With Interest Rates

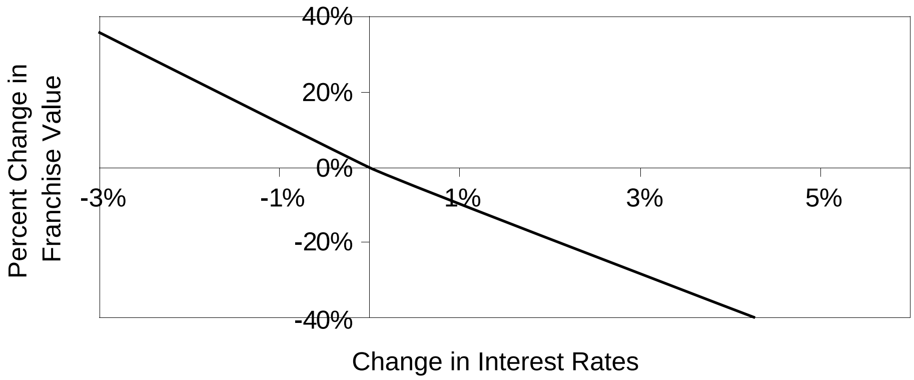

What makes this simplified analysis of franchise value particularly significant is the fact that the insurer can substantially reduce the VaR for its franchise value by altering the duration of its invested assets. In the first case, in which premiums remain constant, the insurer can reduce VaR by extending its asset duration. In the second case, in which the insurer alters premiums or loses business to competitors, it can reduce VaR by shortening the duration of its investments. In either case, by making this change the insurer creates a situation in which the change in its franchise value produced by a change in interest rates is exactly offset by an equal and opposite change in the value of its invested assets. However, by acting appropriately to achieve such a reduction in franchise risk, the insurer will appear to be taking additional risk, since franchise value is not included on its balance sheet. (For a more detailed example and demonstration of these points see Panning 1995.)

This example has several important implications. First, when applying VaR to the firm as a whole, it is important to include in the analysis all components of the firm's value, including those not recognized by current accounting standards. Second, in calculating a firm's VaR for an extended horizon, the analysis should take into account risk factors that are not normally included in short-term VaR calculations, factors such as the effect of competition. This is, in fact, a generalization of the point presented in the above analysis of adaptive risk modification. When VaR is calculated for long-term horizons, it is essential to include in the analysis the likely actions of relevant

Figure 17
Distribution of Changes in Franchise Value
When Premiums Vary Inversely with Interest Rates

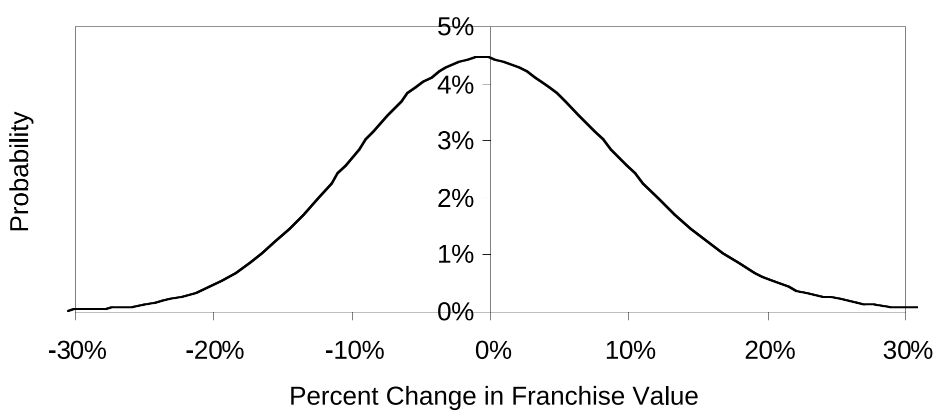

decision makers, including customers and competitors as well as the managers of the firm.

6. CONCLUSION

The concept that underlies the calculation of VaR is in fact quite old. VaR is in principle nothing more than what statisticians call a one-sided confidence interval. What is new is the availability of high-speed computers and historical data needed to implement these calculations for large and complex portfolios of securities. These same resources also make it possible to calculate VaR for assets and liabilities other than securities, for firms other than investment banks, and for decisions that are strategic rather than tactical. But VaR will be useful in these other contexts only if it successfully takes into account the longer time horizons that prevail in them.

As we have seen, time horizons longer than those prevalent in investment banking pose three problems for VaR analysis. Estimation risk—the risk of inaccurately forecasting the response of assets or liabilities to changing conditions—can in fact be quantified and incorporated into VaR by using appropriate statistical methods. Adaptive risk modification—the responses of decision makers to changes in their circumstances—can likewise be incorporated by correctly modeling them and including these responses in each of the scenarios from which VaR is calculated. Finally, franchise risk is appropriately incorporated by refusing to take the firm's balance sheet at face value and by again carefully modeling and including pertinent strategic decisions in the analysis.

The arguments, examples, and solutions presented here by no means exhaust the challenges that must be addressed in extending the scope of VaR to other industries and to decisions that are strategic rather than tactical. Nonetheless, I believe they provide a useful starting point for making greater use of the considerable virtues of VaR.

ACKNOWLEDGMENTS

I am indebted to the Society of Actuaries for its support and encouragement, and to my former colleagues at Willis Corroon and two anonymous referees for helpful comments and criticisms.

REFERENCES

BAWA, V., BROWN, S., AND KLEIN, R. 1979. *Estimation Risk and Optimal Portfolio Choice.* Amsterdam: North-Holland.

BLACK, F., AND SCHOLES, M. 1973. ''The Pricing of Options and Corporate Liabilities,'' *Journal of Political Economy* 81: 637–54.

BUTSIC, R. 1994 ''Solvency Measurement for Property-Liability Risk-Based Capital Applications,'' *Journal of Risk and Insurance* 61:656–90.

CUMMINS, D. 1988. ''Risk-Based Premiums for Insurance Guarantee Funds,'' *Journal of Finance* 43:823–39.

DIXIT, A., AND PINDYCK, R. 1994. *Investment under Uncertainty.* Princeton: Princeton University Press.

DOHERTY, N., AND GARVEN, J. 1986. ''Price Regulation in Property-Liability Insurance: A Contingent Claims Approach,'' *Journal of Finance* 41:1031–50.

JORION, P. 1997. *Value at Risk.* Chicago: Irwin.

MORGAN, J.P. 1995. *RiskMetrics—Technical Document.* New York: J.P. Morgan.

PANNING, W. 1995. ''Asset-Liability Management for a Going Concern,'' in *The Financial Dynamics of the Insurance Industry,* edited by E. Altman and I. Vanderhoof. Burr Ridge, Ill: Irwin, pages 257–91.

SILVERMAN, B. 1986. *Density Estimation for Statistics and Data Analysis.* London: Chapman and Hall.

SPARKS, A., AND SUNG, F. 1995. ''Prepayment Convexity and Duration,'' *Journal of Fixed Income* 4:7–11.

TAYLOR, G. 1987. ''Regression Models in Claims Analysis. I: Theory,'' *Proceedings of the Casualty Actuarial Society* 74: 354–83.

TRIGEORGIS, L. 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation.* Cambridge: MIT Press.

Discussions on this paper can be submitted until October 1, 1999. The author reserves the right to reply to any discussion. Please see the Submission Guidelines for Authors for instructions on the submission of discussions.