Interpreting Value at Risk (VaR)

Allan W. Gregory

Department of Economics, Queen's University, Kingston, Ontario, Canada

awg@qed.econ.queensu.ca

Jonathan J. Reeves

School of Banking and Finance, University of New South Wales, Sydney, Australia

reeves@unsw.edu.au

Abstract

Value at Risk (VaR) has been increasingly accepted globally by both risk managers and regulators as a tool to identify and control exposure to market risk. We introduce a distinction between the structural sources of VaR, and the reduced-form VaR measures that are actually calculated by financial firms. For instance, modern portfolios are characterized by a constantly changing composition of security holdings that reflect portfolio managers' strategies, expected prices, and net cash flows into the portfolio. As a result of these factors, portfolio returns are time-varying mixtures of distributions which are unlikely to be well-approximated by conventional methods. We argue that this gap between the structural sources of VaR and reduced-form measures suggests that these calculations are unreliable and will ultimately fail. We formalize the complex evolution of a portfolio, and relate this to standard VaR calculations. The difficulties with these calculations are illustrated in an empirical example consisting of a portfolio manager who is subject to stochastic net cash flows.

JEL code: G15, G21

Keywords: Financial Risk Management, Value at Risk

1 Introduction

This paper examines conventional methods for calculating Value at Risk (VaR) and argues that the methods are inconsistent with structural factors underlying the actual distribution of portfolio returns. VaR represents the maximal probable loss for a portfolio over a specified trading horizon and stated probability level. Traditional VaR methods, can be viewed as reduced-form measures, that typically assume a time-invariant distribution of portfolio returns, constant security holdings and the absence of net portfolio cash flows over the projected trading horizon. These assumptions neglect important characteristics of the actual portfolio problem, including the role of active trading strategies, unexpected changes in portfolio size and a time-varying distribution of security returns. Reflecting these time-varying structural factors, the distribution of portfolio returns is heterogeneous over time. Further, these factors are unlikely to be well-approximated by conventional time-series models. Although, each of these neglected components may be small relative to the entire portfolio, they will, on occasion, be important enough to lead to persistent violations of VaR especially given the probability levels chosen for these calculations are so very low.

The gap between the structural sources of VaR and the reduced-form VaR measures suggests that these calculations are unreliable and will eventually fail. Despite a dubious connection to any underlying probability level, the reduced-form VaR statistic represents a leading tool employed by risk managers globally to identify and control exposure to market risk. There are several factors underlying its growth as a risk-management tool. First, all manner of firms increasingly hedge their exposure to sources of market risk. Financial market volatility can dramatically impact the profitability of firms, particularly those firms that improperly hedge market risks. The experience of Enron, Orange County, Long-Term Capital Capital Management, Metallgesell Schraft and Barings Bank illustrates the magnitude of large unanticipated events that are not going to be captured by standard VaR calculations. Notwithstanding these outstanding special cases, both private investors and financial-sector regulators require transparent measures of market risk exposure. The VaR approach addresses these issues by offering an appealing summary statistic of portfolio risk embodied within a single statistic. Specifically, VaR is a powerful operational tool to establish position limits for traders, but also enables managers and regulators to control to some extent the firms overall margin between risk and return. The Basle Accord (1988, 1996) requires that financial institutions establish capital reserves at a level closely associated with the VaR characterizing their portfolios.

The formal definition of VaR is the loss $x(t)$ such that the portfolio will lose less than $x(t)$ with probability $(1-\alpha)$ during the forecast horizon. Consistent with this definition, VaR represents the left-tail of a one-sided confidence interval conditional on information available at time t. We denote this statistic as $VaR^{\alpha}_{t+k|t}$. For financial institutions, the forecast horizon $t+k$, probability level α, and the *general* methodology for calculating VaR are dictated by the terms of the Basle Accord (1988, 1996).

There are two widely adopted reduced-form approaches for calculating VaR. First, the variance-covariance approach assumes the joint normality of security returns so that the predicted overall portfolio return is itself normally distributed. Second, the historical simulation method posits that the conditional distribution of future returns resembles the empirical distribution of past returns. For this approach, the VaR is constructed by bootstrapping (independent resampling) repeatedly from the empirical distribution of returns to construct the predicted distribution of future returns. Both of these methods enable rapid computation of VaR for large and heterogeneous portfolios. However, the joint normality approach lacks compelling empirical support and needs modification for portfolios containing nonlinear payoffs. For instance, security returns frequently exhibit various forms of autoregressive conditional heteroscedasticity (ARCH), possible serial correlation and returns that do not straddle the full range of outcomes required under the normal distribution. In addition, the empirical distribution for the historical simulation method may not properly represent future returns if information or market assessment changes. In this case, the data-generating process diverges fundamentally from the past. Further, both methods represent reduced-form measures, so that security positions are held constant over the forecast horizon, returns are governed by a time-invariant distribution, and the portfolio does not incur net cash flows.

In realistic portfolios, our conjecture is that VaR calculations will eventually fail in the sense that the sequence of actual portfolio-loss realizations will not straddle the sequence of left-tail confidence intervals with stated probability α. An assertion to the contrary amounts to a claim that we can accurately predict the conditional distribution of market outcomes. Indeed, conditional forecasts of the mean or median for individual securities have not been highly successful. Thus, it should not be surprising that estimating the conditional tail density for an entire portfolio distribution should prove unattainable. While there have been considerable successes in finance using univariate ARIMA models, vector autoregressive models, and more recently ARCH models, the criterion of success is in-sample-fitting. Calculating accurate tail probabilities is a more difficult task and is likely beyond the reach of current techniques in portfolio analysis.

Despite these concerns, the growth of Value at Risk as a statistical measure of exposure to market risk has led to a variety of new time-series methods to calculate VaR statistics.

The new statistical models are within the conventional reduced-form class and can be broadly categorized into two groups: (i) new statistical representations for the underlying stochastic structure of the portfolio or; (ii) new methods to approximate tail probabilities for a given statistical model. Examples of the former include Engle and Manganelli (2001) who propose a conditional autoregressive VaR technique. Unfortunately, in terms of actual practical applications, these more complicated statistical approaches lend themselves mostly to small and quite artificial test portfolios, which contrasts with the large and heterogeneous portfolios of modern financial institutions. Examples of the latter category focus on extreme value estimation to estimate directly the tails of the distribution.

Most market participants already appreciate these limitations, and interpret the reported VaR statistics of financial institutions accordingly. Notwithstanding its limitations, reduced-form VaR methods convey valuable information, and impose investment discipline for private investors, regulators and portfolio managers. The VaR statistic represents a convenient summary statistic that, in principle, may be compared across portfolios and time, even if the probability interval interpretation is dubious.

In this paper we formalize the complex prediction problem by modelling the portfolio manager's optimization problem, which emphasizes the inherent difficulty in producing reliable confidence intervals for predicted future returns. Against this structural model, VaR calculations can be understood as reduced-form approximations, which by their nature are unlikely to perform well over sufficiently long periods of time. We also broaden the notion of VaR so that it is more readily communicable to senior managers, shareholders and regulators. This refinement to VaR, which we describe as probability-VaR or p-VaR, emphasizes the entire distribution of returns. As an empirical illustration, we evaluate standard VaR in three different environments: (i) fixed weight asset mix, (ii) changing asset mix and (iii) exogenous cash in/outflows, using a portfolio of daily US equity data. The data are compiled over the 1991 to 1999 sample period. The particular equity claims are selected on the basis of published positions held by highly-capitalized US mutual funds.

The paper is organized as follows. Section 2 reviews different extensions to VaR pursued in the recent literature and discusses the inclusion of VaR to the regulatory environment of commercial banking. Section 3 presents a formal model of the portfolio problem, the statistical definition of VaR, and p-VaR. Section 4 contains an empirical evaluation of VaR. Section 5 concludes with brief remarks.

2 VaR in the Literature

Recent interest in Value at Risk has fostered a large body of research organized along three general themes of inquiry. First, several studies have investigated the statistical reliability and theoretical properties of conventional reduced-form techniques. A second branch of the literature posits new VaR approaches that are designed to close the statistical gap between actual exposure to sources of market risk, and the risk indicated by reduced-form VaR measures. Finally, a third extension of the literature investigates the incentive and moral hazard issues confronting commercial banks as a result of regulations imposed under the Basle Accord (1988, 1996).

First, consider those studies evaluating the statistical reliability of conventional reduced-form measures. These studies often find that the calculated VaR for a given portfolio varies significantly with the choice of reduced-form method. This result reinforces our assertion that reduced-form VaR measures are unreliable, and will ultimately fail. In this regard, work by Hendricks (1996), Gizycki and Hereford (1998), Consigli (2002) and Daníelsson (2002), evaluates the statistical reliability associated with different reduced-form VaR techniques. Hendricks (1996) created over 1,000 randomly selected foreign-exchange portfolios involving up to eight currencies over the 1983 to 1994 period. He investigated the calculations generated by commonly used VaR methods at commercial banks within the US. These methods deliver coverage near the 95 percent confidence level, but deteriorate as the interval is widened to the 97.5 and 99 percent levels. The Basle Accord requires that banks implement VaR based on the 99 percent confidence level. The historical simulation approach also yielded larger risk estimates, and was generally more volatile compared to the normality approaches. Consistent with the findings of Kupiec (1995), it is difficult to assess the statistical accuracy of infrequently occurring events because statistical tests lack sufficient power. However, it is clear that the measure of risk exposure is not independent from the choice of reduced-form technique. This finding presents a dilemma for the "pre-commitment" model whereby financial firms select a risk-management model and report the corresponding VaR to regulators.

It has also been found that VaR does not satisfy conditions of a sensible risk measure. For example, VaR is not sub-additive. This means that the VaR of a portfolio can be larger than the sum of the VaRs for the stand-alone components. These violations of the conditions of a sensible risk measure have been discussed most notably by Artzner et al. (1997 and 1999), Acerbi and Tasche (2002), Frey and McNeil (2002) and Szegö (2002).

A second set of studies in the literature posits new approaches to calculating VaR. The historical simulation method draws from the empirical distribution of past returns under the assumption that returns are independently and identically distributed. However, a large body of financial research demonstrates time dependence in financial returns. In the context of VaR measures, Christoffersen, Diebold and Schuermann (1998) and Danielson and de Vries (1997) argue that excluding time dependencies from VaR calculations imposes a significant bias. In response to this problem, a number of authors have proposed VaR methods that incorporate time dependencies, particularly volatility clustering as discussed by Pagan (1996). For instance, Christoffersen (1998) reviews a number of techniques to incorporate time dependence within VaR calculations. Moreover, Engel and Gizycki (1999) evaluate the effectiveness of VaR with respect to the actual portfolios of Australian foreign exchange dealers. Their findings suggest that none of the principle VaR methodologies strongly outperforms any other. Further, the accuracy of the results improved by exponentially-weighting the data, particularly at the 95-percentile. Boudoukh, Richardson and White (1998) proposed this exponentially weighted approach as a means to incorporate short-run time dependencies. Berkowitz and O'Brien (2002) provide direct evidence on the performance of VaR models for six large commercial banks. With daily data spanning from January 1998 to March 2000 they found conservative VaR estimates with the number of times that the actual portfolio loss exceeded the VaR being substantially less than expected based on the probability level.

In addition to this research, a number of studies consider more sophisticated innovations to the VaR approach. For example, Engle and Manganelli (2001) present an estimation method based explicitly on quintile behaviour. In a measure described as CAViaR, the measure of portfolio risk is interpreted as a quintile of possible future returns conditional on current information. A genetic algorithm is applied to optimize the procedure, and allows adaptation of the measure to new risk environments. Application of the procedure remains limited to a restricted portfolio composed of a composite index. Alternative risk measures have also been proposed. Artzner et al. (1999) discuss several such measures, most notably expected shortfall (ES), which measures the expected loss conditional on the VaR being violated. This risk measure has been analyzed in a number of recent papers, see Acerbi and Tasche (2002), Frey and McNeil (2002), Rockafellar and Uryasev (2002), Tasche (2002) and Topaloglou, Vladimirou and Zenios (2002). However, although the ES measure satisfies certain theoretical properties of a desirable risk measure, including the sub-additive property, it is still subject to the same general critique of VaR in that portfolio returns are time-varying mixtures of distributions which are unlikely to be well-approximated by time series methods.

Finally, a third branch of the literature evaluates the behavioural implications that arise from the regulatory framework incorporating VaR methods. Since the advent of the Basle Accord (1988, 1996), regulatory authorities link the reported VaR of financial institutions to the minimum sufficient level of capital reserves. Unfortunately, these regulations have also introduced strategic incentives to misrepresent the reported VaR. These incentives occur because institutions are penalized if they incur violations of the reported VaR exceeding the statistical likelihood of their occurrence. Consequently, the reported VaR statistics are further distorted compared to the actual underlying structural exposure to market risk.

Abstracting from our arguments on the difficulty of actually calculating a VaR at some level α there is a strategic incentive to misreport VaR. The Basle Accord (1988, 1996) represents a departure from traditional mechanisms to regulate the level of capital reserves for financial institutions. The "standard" regulatory approach entails a variable-rate capital charge applied against the different asset classes within the institution's portfolio. For example, a risk-free sovereign bond does not require a corresponding capital reserve. However, a corporate bond that involves default risk requires an 8 percent capital charge to provide sufficient capital in the event of extreme market events. Further, the traditional system of capital charges treats derivatives as high-risk securities despite their function in offsetting risk exposure to underlying assets. Bradley, Wambeke and Whidbee (1991) argue that the 8 percent capital charge under the previous regulatory rules represents an *ad hoc* level, and does not address increasingly volatile financial markets.

The Basle Accord (1988, 1996) replaces this *ad hoc* patchwork of capital charges with a mechanism linking capital reserves to the reported VaR of the institution's portfolio. Specifically, the level of capital reserves is determined as the trailing 60-day average of the calculated VaR multiplied by a scaling factor $f(t)$, where $3 \leq f(t) \leq 4$. The scaling factor incorporates prudence on the part of regulators. The "internal model" approach allows institutions to develop proprietary VaR models given that certain standardized rules are observed. These restrictions include that the VaR is calculated using an $\alpha = .01$ confidence level and the forecast horizon $t+k$ extends 10 days in advance. The institution is free to adopt any approach to calculate their VaR. As Gizycki and Hereford (1998) argue, this VaR approach introduces flexibility for banks to apply their internally developed VaR models within a set of standardized rules. Interestingly, Berger, Herring and Szego (1995) find that the shift to risk-based capital charges may actually increase the level of capital reserves.

The authorities monitor compliance of the VaR framework with a "back-testing" procedure. If a bank exceeds its reported VaR more than 10 times within the 250-day trading year, then the scaling factor increases from three to four. The Basle Accord requires a one percent level of confidence in the VaR calculation, so that a statistically accurate VaR measure will be violated only 2.5 days over the trading year. Back-testing compliance is a feature of the Basle Accord (1988, 1996) and itself introduces strategic and moral hazard incentives to misrepresent the calculated VaR to regulators. For example, if a bank experiences several consecutive violations, which is likely since the volatility of financial returns tends to cluster, then the bank will inflate its reported VaR to over-estimate its risk exposure. Likewise, a firm that experiences few violations may under-report its measured VaR to increase expected returns.

Recent work examines the adverse incentive effects caused by recent VaR regulations. For example, Fusai and Luciano (1998) model Asset and Liability Management (ALM) to illustrate how banks experience an incentive to hide VaR violations to lower capital reserves. In addition, Danielsson, Hartmann and de Vries (1997) introduce a moral hazard problem in an internal markets environment. They demonstrate that the current penalty for excess violations of the reported VaR is too weak to induce desired capital reserves.

3 Portfolio Management and VaR

3.1 The Portfolio Problem

In this section, VaR is defined formally within the context of a general portfolio management problem. At the beginning of period t, the fund manager allocates a given portfolio, $P(t)$, across n different securities indexed by $j = 1, \ldots, n$.

$$P(t) = \sum_{j=1}^{n} \theta_j(t+1) S_j(t) \tag{1}$$

The security prices prevailing at period t for each asset, j, are denoted by, $S_j(t)$. The number of securities purchased in period t of security j is represented by, $\theta_j(t+1)$. Since we want to allow for the introduction of new securities by the market, or the purchase of existing securities at some future date which are currently not in the portfolio, we assume $\theta_j(t+1) \geq 0$ and $S_j(t) \geq 0$. We then interpret $j = 1, \ldots, n$ as indexing all possible securities over the relevant period of analysis.

The time script highlights that shares purchased in period t are carried over into period $t+1$. This is important for calculating portfolio returns, and emphasizes the intertemporal dimension of the portfolio problem.

By construction, the portfolio weights for each security composing the portfolio will sum to one.

$$\delta_j \geq 0 \qquad \text{and} \qquad \sum_{j=1}^{n} \delta_j(t) = 1. \tag{2}$$

The *ex ante* weights are defined as follows,

$$\delta_j(t) = \frac{\theta_j(t+1)S_j(t)}{P(t)}, \tag{3}$$

where the weights are determined in advance of the realization of uncertainty. $\delta_j(t)$ is defined as the current weight of security j in the portfolio using values of the security at time t. In general, the weights will change *ex post* with the realized returns. Consequently, the *ex ante* weights reflect the portfolio manager's risk exposure to each security, j at time t.

At the beginning of the next period $t+1$, the share price $S_j(t+1)$ for each j is realized. As a result, the value of the portfolio changes, and may be expressed as

$$P(t+1) = \sum_j \theta_j(t+1)S_j(t+1).$$

Alternatively, the portfolio at period $t+1$ may be written as,

$$P(t+1) = \sum_{j=1}^{n} \theta_j(t+1)\Big(1 + \frac{S_j(t+1) - S_j(t)}{S_j(t)}\Big) S_j(t), \tag{4}$$

or in terms of security returns between periods t and $t+1$,

$$P(t+1) = \sum_{j=1}^{n} \theta_j(t+1) R_j(t+1|t) S_j(t), \tag{5}$$

where $R_j(t+1|t)$ is the gross return of security j based on the share price at period t. Upon realizing the portfolio value at the beginning of period $t+1$ and using all information currently available, the manager re-optimizes to be consistent with her trading strategy. This strategy is indexed by the following set,

$$\boldsymbol{\theta} = \Big\{\theta_j(t), \quad j = 1, \ldots, n; \quad t = 1, \ldots, T\Big\}, \tag{6}$$

which outlines the strategy over all j securities, and over some finite trading horizon, T. The portfolio manager's strategy will depend on many factors including historical and anticipated returns, expected net new investments in the portfolio, exogenous cash in/outflows, the correlation of returns, current and future purchases, new technologies and sales of the fund, anticipated borrowing conditions and a variety of other financial activities in the market. Such strategies are unlikely to be described by time-invariant decision rules.

The problem of fluctuating net cash flows represents another problematic issue for reduced-form VaR calculations. For instance, many portfolio managers are not certain about net deposits and withdrawals from the portfolio over the forecast horizon of VaR estimates. The random cash flows affect security holdings, and are likely to affect the trading strategy $\{\boldsymbol{\theta}\}$. Random cash flows can be incorporated into this framework by assuming that each financial institution manages two funds. First, the institution passively manages a market-index fund of size $\bar{V}$ (exogenously given). The institution invests the full principal in the market index, but shifts the net return $\bar{V}(R_P(t+1|t)-1)$ into an actively managed fund. The second fund corresponds to the portfolio described above. This assumption calibrates the random cash flows to actual market developments, rather than assuming an exogenous stochastic process to govern net portfolio cash flows. This approach should proxy actual net cash flows since these flows are likely to be correlated with broad market developments. In section four, we investigate the impact on the reliability of reduced-form VaR estimates from introducing net portfolio cash flows calibrated to broad market outcomes.

Define the net cash flow into the actively managed portfolio as $\chi(t+1|t)$. In period $t+1$, the portfolio manager re-balances after the realization of current security prices $S_j(t+1)$, $\forall j$. The portfolio manager experiences the net portfolio cash flow before the re-balancing. As a result, the portfolio value at $t+1$ before re-balancing can be expressed as,

$$P(t+1) \;=\; \sum_{j=1}^{n} \theta_j(t+1)S_j(t+1) \;+\; \chi(t+1|t)\ , \tag{7}$$

3.2 Value at Risk

VaR is defined as a dollar amount, $x(t)$, that the portfolio losses will be less than $x(t)$, with probability $(1-\alpha)$ over a specified time interval (say, k). This notion may be described more formally as,

$$VaR^{(\alpha)}(t+k|t) = x(t), \tag{8}$$

such that,

$$Pr\Big\{(P(t) - P(t+k))|t \leq x(t)\Big\} = 1 - \alpha.$$

Note that $x(t)$ is positive, and is defined in terms of losses. The VaR is usually calculated for a single value of α (say .01). From an investor's perspective of assessing risk, without knowing the current value of the portfolio $P(t)$, the VaR in equation (8), may not be all that informative even given all our earlier arguments. For this reason, we propose a probability-based VaR measure, p-VaR,

$$p - VaR^{\delta}(t+k|t) = Pr\Big\{P(t+k) \leq (1-\delta)P(t)\Big\} \qquad \alpha \geq 0. \tag{9}$$

For instance a p-VaR$^{.05}(t+k|t) = .03$ would indicate that, based on historical simulations, there is a 3% chance that the portfolio will lose at least 5% of its value. Notice that the p-VaR is free of currency units, and may be easily understood as it is analogous to a standard probability value (p-value). In this way, one could trace out the p-VaR, based on historical simulations for the relevant range of δ. More importantly graphs of such information would provide investors a more complete picture and allow them to apply their own risk standards in making decisions and is of course, exactly the intent in classical hypothesis tests of p-values. Another advantage of the p-VaR is that it facilitates a comparison of the entire distribution of losses of the portfolio returns. In contrast, conventional VaR typically analyzes only the a single point of the lower tail of the distribution.

Figure 1 displays the cumulative distribution function based on historical simulation for a portfolio on 21 June 1999. This forecast is carried out ten trading days beforehand based on the assumption that the portfolio holds one millionth of the number of shares outstanding in nineteen popular US stocks. These stocks are listed in section four. The corresponding p-VaR is displayed in Figure 2.

3.3 Calculating VaR: Conventional Methods

There are now a number of different approaches to calculate VaR. Certainly, as we have indicated, many of the more complicated methods are not practical or feasible for modern portfolios comprising hundreds of securities. We restrict discussion to three conventional methodologies to calculate VaR. An excellent exposition on these methods may be found in Jorion (2001), and Linsmeiser and Pearson (1996). The three predominant methods are; (i) historical simulation; (ii) variance-covariance based on normality; and (iii) Monte Carlo simulation from a known distribution. We briefly discuss each in turn.

The first method is based on historical simulation. The basis of this method is to construct a distribution of potential returns based on the empirical distribution of historical returns over the previous N periods. Possible future values for the portfolio are calculated assuming that realized historical returns represent the distribution of future returns. This method holds the number of shares $\theta_j(t+1)$ constant over the forecast horizon k and uses *iid* draws from past returns $R_j(s)$ $s = 1, \ldots, t$ to construct artificial portfolios $P^r(t+i)$, $r = 1, \ldots, R$ and $i = 1, \ldots, k$. This yields R possible future portfolios $P^r(t+k)$ which are ordered form smallest to largest so that $x(t)$ is easily calculated depending on the chosen α. At the very best, this approach only mimics the uncertainty in the price of securities and does not capture the changing composition or mix of assets in the portfolio or net cash flows in the portfolio.

In figures 1 and 2 we have employed historical simulation, however, in these cases instead of sampling 1-period changes repeated to build up the 10-period return, the re-sampling is on 10-period returns directly. Formally, such a sampling is identical to block bootstrapping with a fixed window width of 10, and might produce a better approximation to volatility clustering observed in many financial time series.

A second method used to determine VaR is the "variance-covariance" (normality) approach. This approach is employed by RiskMetrics, the leading risk management consulting firm. This method also fixes θ_j. A critical assumption underlying this method regards the joint behaviour of the underlying risk factors and individual securities. Returns are assumed to be time-invariant serially uncorrelated, and distributed as multivariate normal. The method maps the standard deviation and correlation of the underlying risk factors into an overall portfolio return. The mean and standard deviation matrix of the factor returns are estimated from the historical data. For instance, RiskMetrics is well known for its large correlation matrix of asset returns in numerous currencies that it publishes over the Internet. An advantage of this method is that the benefits of hedging due to securities with negatively correlated returns may be incorporated. The overall portfolio standard deviation can then be derived, which allows for a simple VaR calculation.

An example of the variance-covariance approach is helpful. Let $\boldsymbol{\gamma}$ represent a vector of security sensitivities to the underlying risk factors. Let $\hat{\boldsymbol{\mu}}$ be the estimated mean change for the portfolio and $\widehat{\boldsymbol{\Sigma}}$ be the estimated variance-covariance matrix of the 1-period portfolio change conditional on information available during period t. The change in portfolio value, defined as $\mathbf{\Delta}P|t$, is assumed to follow,

$$\mathbf{\Delta}P|t \sim N(\hat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\Sigma}}). \tag{11}$$

The corresponding VaR is,

$$VaR^{\alpha}(t+k|t) = -k\boldsymbol{\gamma}'\hat{\boldsymbol{\mu}} - Z(\alpha)(k\boldsymbol{\gamma}'\widehat{\boldsymbol{\Sigma}}\boldsymbol{\gamma})^{\frac{1}{2}}, \tag{12}$$

where $Z(\alpha)$ is the α percentile of the standard normal distribution. Hence, the VaR for this portfolio varies proportionally with the volatility in the underlying risk factors.

The third dominant approach for calculating VaR is by Monte Carlo simulation. This method is similar to the historical simulation approach, except that the risk manager *selects* a distribution of potential returns. A pseudo-random number generator draws from the selected distribution B times, where B is a large number typically exceeding 10,000. Subsequently, the distribution of portfolio returns is constructed. From this distribution of portfolio values, the VaR is determined at the α percent level of confidence.

There are strengths and weaknesses associated with each of these general approaches. The historical simulation methods are sensitive to the particular sample of data. For instance, the sample may be characterized by historically abnormal volatility or other price effects. The resulting VaR would reflect this volatility bias, and provide a misleading measure of risk exposure. To some extent, this problem is endemic because the calibrated values in the variance-covariance approach are estimated using the same historical data. Consequently, the variance-covariance method is also dependant on the time-series properties to the selected sample period. In particular, the variance-covariance method may be vulnerable to outliers.

The variance-covariance and Monte Carlo approaches are also problematic because of distributional assumptions. It has been well documented that asset returns exhibit "fat tails" relative to the multivariate normal distribution. The bias may be very important because the VaR statistic focuses primarily on the lower tail of a distribution. Moreover, securities with nonlinear payoff structures are inconsistent with normally-distributed returns. There are modifications to these methods to handle option prices and other derivatives with nonlinear payoff schedules. Although such complications present another complicating twist to accurate calculation of the VaR, we will not pursue them here.

Our more fundamental question than the distribution of $P(t)$, has to do with the statistical underpinning of the VaR calculation itself. By failing to model the underlying sources of risk, raises serious doubts about any VaR approach. Empirical work in the following section with actual portfolios confirms this fact. Monte Carlo studies comparing alternative methods by assuming some time-invariant data generating process are too artificial to be convincing in actual portfolio situations. On the other hand, actual calculations that suggest "good properties" (ie. roughly the "right" x_t for the chosen α) may simply reflect good fortune, since the underlying statistical foundation for any VaR calculation is tenuous.

4 Empirical Performance of VaR

In this section we examine the empirical performance of VaR for a representative fund manager. Our fund manager holds nineteen U.S. stocks in her portfolio in proportion to the actual market weights. These securities include nineteen of the most widely held stocks by US mutual funds; Cisco Systems, Microsoft, Intel, Texas Instruments, Sprint, Amgen, Interpublic Group, Mellon Bank, Warner Lambert, Bristol Myers Squibb, Enron, General Electric, Time Warner, Exxon, Dell Computer, American Express, Sun Microsystems, Corning and Ford Motor. Trading data for the security returns between 1991 and 1999 are obtained from the CRSP database.

The 95% and 99% 10-day-ahead VaR based on normality (*VCov*) and historical simulations (*Historical*) are computed for each trading day. The previous 500 days of return data are used to calculate the VaR statistic. For the historical simulation approach, VaR is calculated by drawing from the empirical distribution based on the past 500 observed security returns. For the variance-covariance approach, the volatility and covariance parameters are calibrated from the the same 500-day sample of return data. The weights of each security used for the calculation of VaR are based on the actual market weights prevailing on the day of trading in the sample. Our representative fund manager holds one millionth of the number of shares outstanding in each of these stocks.

We evaluate the effectiveness of these reduced-form VaR calculations using a standard back-testing procedure applied in practice by the policy authorities and commercial bankers. Specifically, we investigate if the number of portfolio values that departed from the predicted VaR level exceeds the confidence level. The results of this back-testing exercise are reported in Tables 1 and 2. Several experiments are considered in this empirical back-testing exercise using both the *VCov* and *Historical* approaches in each case. The first experiment holds the number of securities fixed over the VaR forecast horizon. This case is denoted as *Fix* in the tables of reported results. The second experiment introduces a change in the portfolio mix of security holdings during the 10-day forecast horizon. Standard calculations assume fixed portfolio holdings during the forecast period.

This case is denoted as *Mix* in the tables. Finally, we consider the realistic feature that portfolio size is subject to random inflows or outflows. The experiment is labelled under *Cash* in the tables of results.

First, consider the empirical findings for the *Fix* experiment. Over each ten-day period, the realized return is computed based on the number of shares held constant over the 10 day period, although the data are adjusted for stock splits, rights issues and other distributions or disbursements. Under the *Fix* category the number of times over a year that this loss was greater than the corresponding VaR is calculated. An accurate VaR estimate would have approximately 2-3 violations of the 99% VaR for each year. In contrast, we can observe from the tables that there are four years without any violations under either the *VCov* or *Historical* procedure for calculating VaR. Moreover, there are only one and two violations in each of the remaining two years, respectively. At the 95% confidence level, we would expect an accurate procedure to produce 12-13 violations of the VaR per year. The empirical results demonstrate significant under-reporting of violations, which suggests that VaR overestimates market risk for any given confidence level.

Second, consider the *Mix* category. This experiment emphasizes how active trading strategies affects the relationship between actual portfolio returns during the forecast horizon and the calculated VaR. One active trading strategy that we consider involves the representative fund manager selling the "dogs" and reinvesting the proceeds into the "winners" on the fifth day of the 10-day forecast horizon. Specifically, the fund manager sells the nine stocks with the lowest holdings by US mutual funds and reinvests the proceeds equally among the 10 remaining stocks in the portfolio. The results reported in the tables suggest that the number of VaR violations substantially increase during some years, most notably 1994 and 1996. The largest increases at the 99% confidence level occur in 1994 with the violations increasing from zero to three. At the 95% confidence level, we observe violations in 1994 increase from 8 to 19, under both the *VCov* and *Historical* methods.

Finally, we investigate the case of a cash inflow or outflow into our representative fund manager's portfolio. The cash flow change is calibrated to the corresponding 10-day return on the CRSP value-weighted index, which includes dividends. The percentage change in this index over the VaR forecast period is applied to the original market holdings in our portfolio to produce a new value for the portfolio, and thus a new realized return over this time period. The number of portfolio returns that violate the predicted VaR are displayed in Tables 1 and 2 under the *Cash* heading. The results illustrate significant changes in the number of violations in the VaR level. Specifically, consider the results for 1997. For the *Fix* experiment, there are two violations at the 99% level. However, under the *Cash* approach, there are 15 and 14 violations using the *VCov* and *Historical* methods, respectively. This compares to only two or three violations if the VaR is accurately measured at the 99% level.

These empirical findings demonstrate that standard VaR estimates are not accurate for measuring the lower percentiles of the loss distribution of a portfolio, even when the security holdings are held fixed over the period. When other realistic features of the evolution of a portfolio are incorporated, such as a changing mix in security holdings or varying cash flows, we observe significant changes in the number of violations compared to the predicted VaR level. Thus, we have demonstrated that reduced-form methods to calculate VaR can substantially under-estimate the actual exposure of a portfolio to market risk.

5 Conclusions

Value at Risk has become the standard measure of financial risk employed worldwide. However, we have shown that the statistical estimation of VaR is often flawed as the uncertainties in which a VaR estimate is attempting to approximate are far too complex for standard time series techniques. We have demonstrated this by firstly formalizing the evolution of a portfolio and comparing this process to standard VaR measures. Second, we have empirically demonstrated that the actual losses that a portfolio can experience, can be far greater and more variable than what a VaR measure suggests. This undoubtedly leads to firms' VaR estimates providing an inaccurate evaluation of the actual market risks in which the firm is exposed. We feel that the current VaR estimation techniques are unreliable as a primary measure of risk, which suggests that additional work is required to develop more reliable measures.

Acknowledgments

The authors would like to thank Todd Mattina for extensive comments and suggestions and are also grateful to Thomas Henker and Derek White for assistance in data collection. We acknowledge financial support from the Social Sciences and Humanities Research Council of Canada.

Appendix

Table 1: Number of Violations of the 10 Day VaR at $\alpha = 0.01$

Year	*VCov*			*Historical*		
	Fix	Mix	Cash	Fix	Mix	Cash
1994	0	3	3	0	3	3
1995	0	1	0	0	0	0
1996	1	3	8	1	3	5
1997	2	2	15	2	2	14
1998	0	0	2	0	0	2
1999	0	0	0	0	0	0

Table 2: Number of Violations of the 10 Day VaR at $\alpha = 0.05$

Year	*VCov*			*Historical*		
	Fix	Mix	Cash	Fix	Mix	Cash
1994	8	19	13	8	19	13
1995	11	15	1	11	15	1
1996	10	18	23	10	18	22
1997	11	13	20	10	13	21
1998	2	3	14	2	3	15
1999	2	1	2	2	1	2

Note: *VCov* is the variance covariance calculation based on normality and *Historical* is historical simulations with 10000 simulations employed. *Fix* refers to the case where the number of shares of each security is held constant at the initial value. *Mix* allows the security mix to change by the sell off of the nine least popular stocks after 5 days with the funds re-invested equally over the remaining ten stocks. *Cash* allows the portfolio to have net inflows and outflows of cash according to the 10-day return on the CRSP value-weighted index (including dividends).

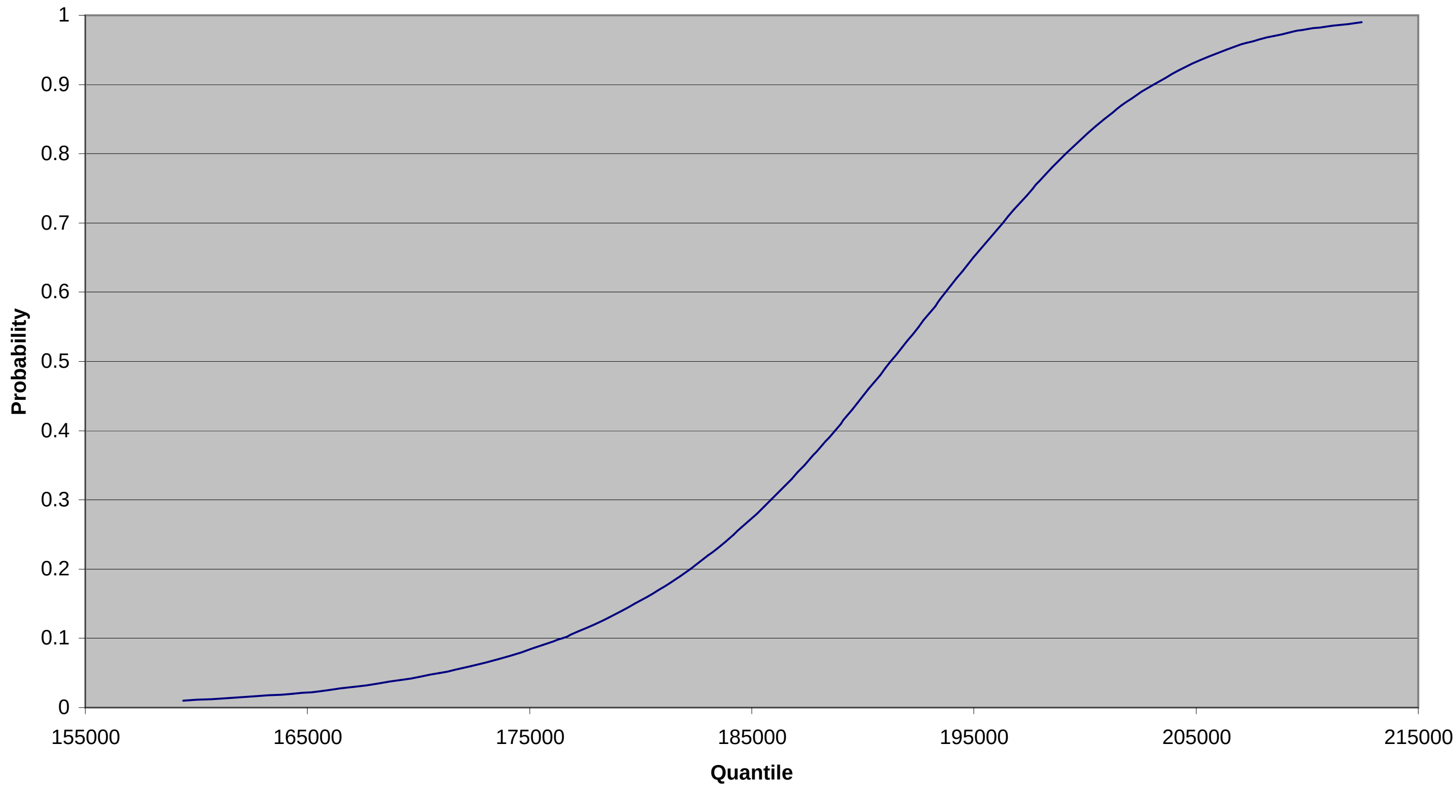

Figure 1: Cumulative Distribution Function (Historical Simulation) for Portfolio on 21 June 1999

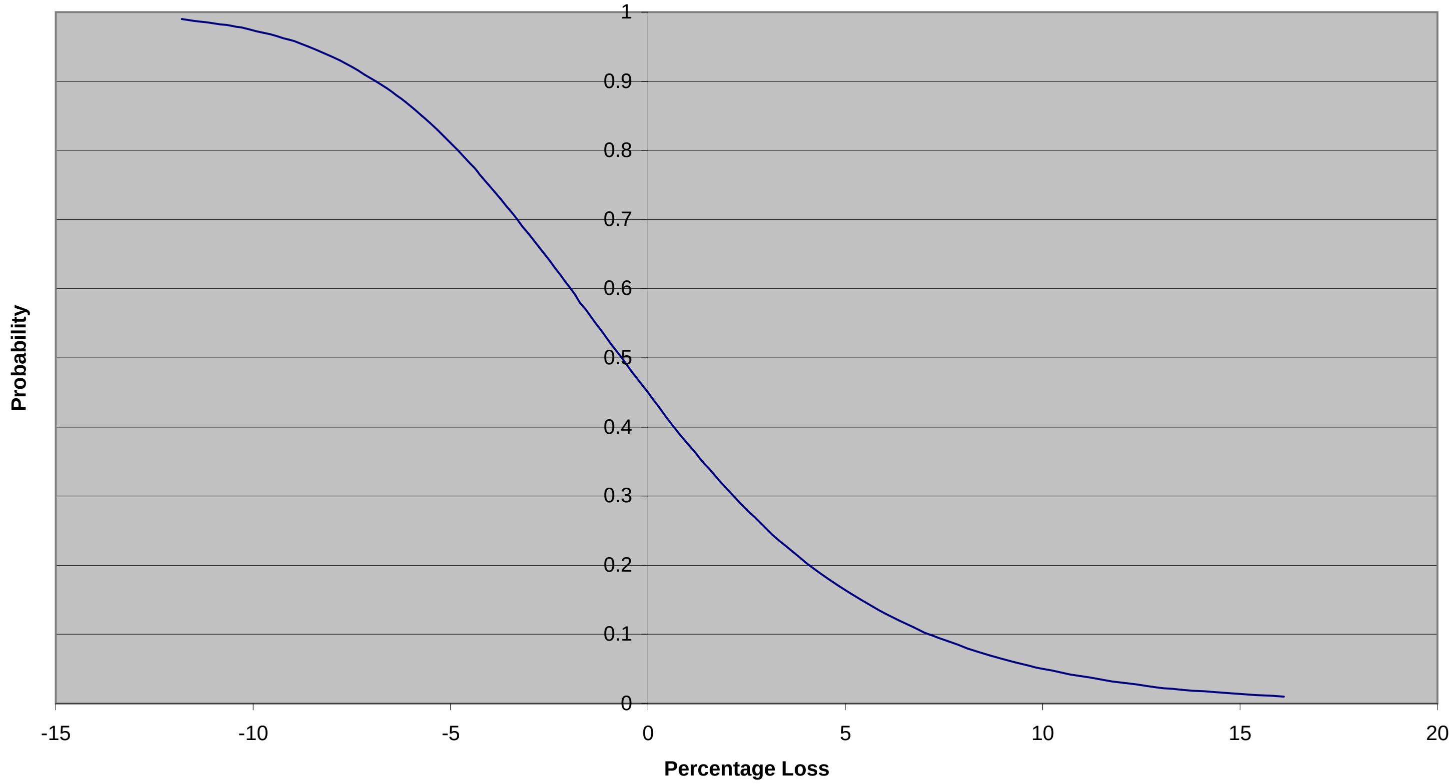

Figure 2: P-VaR(Historical Simulation) for Portfolio on 21 June 1999

References

Acerbi C., Tasche, D., 2002. On the coherence of expected shortfall. *Journal of Banking and Finance*, 26, 1487-1503.

Artzner, P., Delbaen, F., Eber, J-M., Heath, D., 1997. Thinking coherently. *Risk*, 10 (11).

Artzner, P., Delbaen, F., Eber, J-M., Heath, D., 1999. Coherent measures of risk. *Mathematical Finance*, 9, (3) 203-228.

Basle Committee on Banking Supervision, 1988. International Convergence of Capital Measurements and Capital Standards. *BIS,* Basle, July.

Basle Committee on Banking Supervision, 1996a. Supervisory Framework for the use of 'Backtesting' in Conjunction with the Internal Models Approach to Market Risk Capital Requirements. *BIS,* Basle, Report No. 22.

Basle Committee on Banking Supervision, 1996b. Amendment to the Capital Accord to Incorporate Market Risks. *BIS,* Basle, Report No. 24.

Berger, A., Herring R., Szegö, G., 1995. The role of capital in financial institutions. *Journal of Banking and Finance*, 19, 393-430.

Berkowitz, J., O'Brien, J., 2002. How accurate are value-at-risk models at commercial banks? *Journal of Finance*, 57, (3) 1093-1111.

Boudoukh, J., Richardson, M., White R., 1998. The best of both worlds. *Risk*, 11, 64-67.

Bradley, M., Wambeke, C., Whidbee, D., 1991. Risk weights. risk-based capital and deposit insurance, *Journal of Banking and Finance*, 15, 825-893.

Consigli, G., 2002. Tail estimation and mean-VaR portfolio selection in markets subject to financial instability. *Journal of Banking and Finance*, 26, 1355-1382.

Christoffersen, P., 1998. Evaluating interval forecasts. *International Economic Review*, 39(4), 841-862.

Christoffersen P., Diebold F., Schuermann, T., 1998. Horizon problems and extreme events in financial risk management. *Federal Reserve Bank of New York Economic Policy Review*, 4(3), 109-118.

Daníelsson, J., 2002. The emperor has no clothes: Limits to risk modelling. *Journal of Banking and Finance*, 26, 1273-1296.

Daníelsson, J., Hartmann, P., de Vries, C., 1997. The cost of conservativeness: Extreme returns, value-at-risk, and the Basle multiplication factor. University of Iceland Working Paper.

Engel, J., Gizycki, M., 1999. Conservatism, accuracy and efficiency: Comparing value-at-risk models. *Discussion Paper 2*, Australian Prudential Regulatory Authority.

Engle, R., Manganelli, S., 1999. CAViaR: Conditional autoregressive value at risk by regression quantiles. *University of California, San Diego, Department of Economics Working Paper*, 99(15).

Frey, R., McNeil, A., 2002. VaR and expected shortfall in portfolios of dependent credit risks: Conceptual and practical insights. *Journal of Banking and Finance*, 26, 1317-1334.

Fusai, G., Luciano, E., 1998. Measuring VaR under optimal and suboptimal portfolio policies. Manuscript, University of Torino.

Gizycki, M., Hereford, N., 1998. Assessing the dispersion in banks' estimates for market risk: The results of a value-at-risk survey. *Discussion Paper 1*, Australian Prudential Regulatory Authority.

Hendricks, D., 1996. Evaluation of value-at-risk models using historical data. *Federal Reserve Bank of New York Economic Policy Review*, 2(1), 39-69.

Hendricks, D., Hirtle, B., 1997. Bank capital requirements for market risk: The internal models approach. *Federal Reserve Bank of New York Economic Policy Review*, 3(4), 1-12.

Kupiec, P., 1995. Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3, (Winter), 73-84.

Jorion, P., 2001. Value at Risk: The New Benchmark for Managing Financial Risk. McGraw-Hill, New York.

Lucas, A., 2001. Evaluating the Basle guidelines for backtesting banks' internal risk management models. *Journal of Money, Credit and Banking*, 33(3), 826-846.

Pagan, A., 1996. The econometrics of financial markets. *Journal of Empirical Finance*, 3, 15-102.

Rockafellar, R., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26, 1443-1471.

Szegö, G., 2002. Measures of risk. *Journal of Banking and Finance*, 26, 1253-1272.

Tasche, D., 2002. Expected shortfall and beyond. *Journal of Banking and Finance*, 26, 1519-1533.

Topaloglou, N., Vladimirou, H., Zenios, S., 2002. CVaR models with selective hedging for international asset allocation. *Journal of Banking and Finance*, 26, 1535-1561.