Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?

Xiogwei Ju and Neil D. Pearson

Xiogwei Ju
Ph.D. Candidate
Department of Finance
University of Illinois at Urbana-Champaign

Neil D. Pearson
Associate Professor
Department of Finance
University of Illinois at Urbana-Champaign

OFOR Paper Number 98-08
October1998

Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?

Xiongwei Ju[*†] Neil D. Pearson[‡]

University of Illinois at Urbana-Champaign

September 1998

[*] We thank Wayne Ferson for a helpful discussion, and Zvi Wiener for finding an error in an earlier version. Xiongwei Ju was supported by the Corzine Assistantship of the Office for Futures and Options Research of the University of Illinois at Urbana-Champaign.

[†] Ph.D. candidate, Department of Finance, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, Illinois 61820. E-mail: xju@uiuc.edu; voice: 217 244 2239; fax: 217 244 3102.

[‡] Associate Professor, Department of Finance, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, Illinois 61820. E-mail: pearson2@uiuc.edu; voice: 217 244 0490; fax: 214 244 9867.

Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?

Abstract

We study a source of bias in value-at-risk estimates that has not previously been recognized. Because value-at-risk estimates are based on past data, a trader will often have a good understanding of the errors in the value-at-risk estimate, and it will be possible for her to choose portfolios for which she knows that the value-at-risk estimate is less than the "true" value at risk. Thus, the trader will be able to take on more market risk than risk limits based on value-at-risk permit. Biases can also arise if she does not have a good understanding of the errors, but uses the estimated covariance matrix to achieve certain portfolio objectives. We assess the magnitude of these biases for three different assumptions about the motivations and behavior of the trader and find that in all cases, value-at-risk estimates are systematically downward biased. In some circumstances the biases can be very large. Our study of the distributions of the biases also suggests a way to adjust the estimates to "correct" the biases.

Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?

The risk measurement technique known as value-at-risk has recently become a standard approach for measuring the market risk of financial and commodity derivative instruments, and other financial instruments. Value-at-risk models provide a probabilistic measure of the "market risk" of a portfolio of financial instruments, i.e. the risk that the market value of a portfolio of financial instruments will change as a result of changes in interest rates, foreign currency exchange rates, commodity prices, or equity prices. Specifically, value-at-risk models measure the loss that will be exceeded with a specified probability over a specified time horizon. For example, if the specified probability is 5 percent and the time horizon is one day, then a value-at-risk of $1 million means that the daily mark-to-market loss will exceed $1 million with a probability of only 5 percent. These models have recently become popular because market risk is a key concern of companies' senior managers, investors, and regulators, and the models aggregate the several components of market risk into a single summary measure.

Value-at-risk is used for controlling traders and risk management staff (e.g., setting position and trading limits), determination of capital requirements, performance evaluation, and disclosure to both internal (senior management and/or the board of directors) and external constituencies (regulators and investors). Currently, value-at-risk is regarded as "best practice" for market risk measurement by derivatives dealers and other financial institutions. In addition, it is increasingly used by non-financial corporations, and has recently attracted the interest of regulators. For example, its use is strongly encouraged by banking regulators, and it is one of the three permitted disclosure alternatives in the SEC's recent rule requiring that corporations prepare and disclose quantitative measures of the market risks of their financial instruments. Linsmeier and Pearson (1997b) describe the SEC's new rule requiring disclosure of quantitative measures of market risk.

The three basic approaches for measuring value-at-risk are termed historical simulation, the delta-normal (or analytic or variance-covariance) method, and Monte Carlo simulation. These basic methods are described in Linsmeier and Pearson (1996, 1997a). Zangari (1996) describes the delta-gamma approach, which extends the delta-normal method to instruments with non-linear value functions. Pritzker (1996) and Robinson (1996) describe a number of variants of the basic methods, focusing on the tradeoff between speed and accuracy. Butler and Schachter (1997) and Danielson and de Vries (1997) suggest using kernel estimators in conjunction with the historical simulation method. In addition, Butler and Schachter (1997) and Jorion (1997) provide measures of the precision of the value-at-risk estimate.

All of the methods for computing value-at-risk involve various approximations and estimates, and a number of limitations of both value-at-risk and the methodologies for computing value-at-risk estimates are well understood. All methods assume that the portfolio is fixed over the time horizon used in the value-at-risk calculation, which is usually not the case. The delta-normal method is based on a linearization of the portfolio, and thus can perform poorly with portfolios that include large positions in options or instruments with option-like payoffs (Guldimann (1994)). This is documented by Beder (1995), Jordan and McKay (1995), and Pritzker (1996). In addition, Marshall and Siegel (1997) document the existence of "implementation risk," in that different value-at-risk software will yield different results even while using the same methodology (J.P. Morgan's RiskMetrics™) and data (the standard RiskMetrics™ dataset).

At least equally importantly, value-at-risk estimates are estimates of market risk, based on past data. Mahoney (1995) and Hendricks (1996) provide evidence on the performance of different methods for computing value-at-risk. Alexander (1996), Alexander and Leigh (1997), and Boudoukh, Richardson, and Smith (1997) study the methods used to estimate the variances and covariances used in value-at-risk calculations. Longerstaey (1996) and Duffie and Pan (1997) discuss a range of statistical issues that arise in the estimation of value-at-risk. In addition, Kupiec (1995), Lopez (1997) and Crnkovic and Drachman (1996) discuss statistical methods for evaluating value-at-risk models.

In this paper we focus on the delta-normal method, and study a source of bias in value-at-risk estimates that has not previously been recognized. Specifically, because value-at-risk estimates are based on past data, on any day the trader or trading desk (or other persons who decide which instruments to buy or sell) is likely to have a good understanding of the errors in the value-at-risk estimate. For example, she is likely to know for which markets and instruments historical estimates of market volatility underestimate current market volatility, and for which markets and instruments historical estimates overestimate current market volatility. She is also likely to have information about the relation between current market correlations and historical estimates of them. As a result, it will be possible for her to choose portfolios for which she knows that the value-at-risk estimate is less than the "true" value at risk, and thereby take on more risk than risk limits and/or her supervisor permit. To the extent that she does this, the estimated value-at-risk will be downward biased, i.e. the "true" value-at-risk will exceed the estimated value at risk.

Furthermore, the value-at-risk estimate can be biased even if the trader relies on the estimated market variances and covariances and does not have knowledge of the "true" covariance matrix. If she uses the estimated covariance matrix to achieve certain trading or hedging objectives, and also computes the value-at-risk using the same estimated covariance matrix, on average she underestimates the risk, possibly to a large extent. For example, suppose she determines a hedge based on the estimated covariance matrix of two assets, and then, after establishing the hedge, she estimates the risk of her hedged portfolio using the same estimated covariance matrix. In this case she is underestimating the risk. Objectives that can lead to biases include maximizing expected return subject to a constraint on the estimated value-at-risk, minimizing portfolio standard deviation subject to a constraint on expected return, and certain other objectives involving the estimated variance-covariance matrix. These biases are caused by the sampling error in the estimated variance covariance matrix, which follows a Wishart distribution. If the trader selects uses the sampling error to select a portfolio to achieve a small estimated portfolio standard deviation, then the risk will be underestimated. In the cases we discuss in this paper, the trader is likely to choose such portfolio weights.

Specifically, we assess the magnitude of these biases in value-at-risk estimates for three different assumptions about the motivations and behavior of the trader. We first consider a trader who seeks to maximize "true" value-at-risk subject to a constraint on estimated value-at-risk. This corresponds to a situation in which the trader is seeking to evade risk limits, and addresses the question "how wrong can you be?" While this may seem like an extreme case, it is relevant because value-at-risk has been suggested for use in the "control" function. In this context, it is reasonable to consider a trader who is trying to evade risk limits - preventing this is one of the main objects of the "control" function. Second, we consider a more typical case of a trader who seeks to maximize expected return subject to a constraint on estimated value-at-risk. Finally, we assume that a trader has identified a preferred portfolio, but is unable to hold it because the estimated value-at-risk of the portfolio exceeds some specified limit. In this case we assume that the trader seeks to hold a portfolio as close as possible to the preferred portfolio, subject to the constraint on estimated value-at-risk. For each of these assumptions, we determine the bias for different assumptions about the number of different instruments to which she has access[1] and the number of observations used in estimating the covariance matrix. In all cases, we are able to show that the

[1] For example, whether she has access to U.S. dollar denominated fixed income instruments, U.S. dollar and U.K. pound denominated fixed income instruments, fixed income instruments in all of the actively traded currencies, etc.

distribution of the bias does not depend on the "true" covariance matrix generating the data. In addition to simplifying the presentation of the results, this allows determination of the distribution of the bias without knowledge of the true covariance matrix, making feasible adjustment of the estimates to "correct" the bias.

In the case of a trader who seeks to evade risk limits and take on as much risk as possible, the bias is large except when the number of available assets is small (i.e., less than or equal to 20) and the number of observations used in estimating the covariance matrix is large (greater than or equal to 500). In the other two cases, the bias in estimated value-at-risk is smaller, but still large for some reasonable combinations of parameters. In particular, the bias is very large when we estimate the covariance matrices by weighting the data using exponentially declining weights. This raises concerns about the use of this approach.

Our results apply to the use of value-at-risk in the control and performance evaluation of an individual decision making unit such as a trader or trading desk. They also apply to companies (e.g., some corporate end-users of derivatives) in which the entire portfolio of debt, derivatives, and other financial instruments is centrally controlled by a single decision making unit. For firms or companies with multiple trading desks, whether and how the biases at the level of the individual decision making unit aggregate to biases at the firm level will depend upon the correlations among the portfolios chosen by the individual units.

In the next section we briefly describe the set-up, and distinguish between estimated and "true" value-at-risk. Then in Section II we consider a trader who seeks to maximize "true" value-at-risk subject to a constraint on estimated value-at-risk. This is the case where the bias is greatest. In Sections III and IV we consider traders who either seek to maximize expected return subject to a constraint on estimated value-at-risk, or seek to hold a portfolio as close as possible to a preferred portfolio, subject to a constraint on estimated value-at-risk. Section V briefly concludes. Proofs of our claims are in the appendix.

I. Estimated and "True" Value-at-Risk

In order to focus on the potential biases, we consider the computation of delta-normal value-at-risk in a very simple set-up. Specifically, there are K assets, with prices on the n-th date denoted by the $K \times 1$ vector p_n. These K assets may be interpreted as the "standardized positions" often used in value-at-risk systems (see, for example, Guldimann 1994). The (absolute) price changes $x_n = p_n - p_{n-1}$ are assumed to be draws from a multivariate Normal distribution with a mean vector m and a (non-singular) covariance matrix Σ. The portfolio held by the trader or trading desk is represented by a $K \times 1$ vector w giving the units (e.g., number of Japanese yen, not the fraction of wealth invested in Japanese yen) of each of the K assets held by the trader. We are interested in the risk of various portfolios w.

The risk manager does not know Σ, but rather possesses only an estimate

$$\hat{\Sigma} = \sum_{n=1}^{N} l_n (x_n - m)(x_n - m)' \tag{1}$$

constructed using N observations and a set of weights $\{ l_1,\ l_2,\ l_3, \ldots,\ l_N \}$, where $\sum_{n=1}^{N} l_n = 1$.

Equation (1) includes as special cases both the equally weighted covariance matrix estimator

$$\hat{\Sigma} = \tfrac{1}{N} \sum_{n=1}^{N} (x_n - m)(x_n - m)' \tag{2}$$

and the exponentially weighted estimator

$$\hat{\Sigma} = (1 - l) \sum_{n=1}^{\infty} l^{n-1} (x_n - m)(x_n - m)' \tag{3}$$

or

$$\hat{\Sigma} = (1 - l) \sum_{n=1}^{N} l^{n-1} (x_n - m)(x_n - m)', \tag{4}$$

where in equations (3) and (4) $l < 1$ and N is chosen to be large enough so that the omitted terms have a negligible impact on the sum. In practice, in estimating the covariance matrix it is commonly assumed that $m = 0$, because for the data commonly used in financial applications the mean has only trivial impact on the estimate of the covariance matrix (see, e.g. Figlewski (1997)). We make this assumption below in analyzing the estimated value-at-risk.

Using $\hat{\Sigma}$, the estimate of the portfolio variance is $w' \hat{\Sigma}_N w$ and the estimated value-at-risk is

$$\text{estimated value-at-risk} = k\sqrt{w'\hat{\Sigma}w}\,,$$

where k is a constant, determined by the probability level of the value-at-risk estimate (often k = 1.645 or 2.326). In contrast, the “true” value at risk is

$$\text{"true" value-at-risk} = k\sqrt{w'\Sigma w}\,,$$

where Σ is the actual, in contrast to the estimated, covariance matrix of changes in the market values of the positions. We focus on the extent to which the estimated value-at-risk $k\sqrt{w'\hat{\Sigma}w}$ provides a biased estimate of the "true" value-at-risk $k\sqrt{w'\Sigma w}$.

The potential bias originates in the fact that for many value-at-risk systems $K > N$, i.e. the dimension of the covariance matrix exceeds the number of observations used to estimate it. In many actual value-at-risk systems, K, the dimension of Σ and $\hat{\Sigma}$, exceeds 400. However, the covariance matrix may be estimated with fewer than 400 observations, and as a result the estimated covariance matrix $\hat{\Sigma}$ is singular. This implies that there are many risky portfolios for which the estimated portfolio variance $w'\hat{\Sigma}w$, and therefore the estimated value-at-risk $k\sqrt{w'\hat{\Sigma}w}$, are zero. Were a trader permitted to execute trades in all markets, it would be possible for her to enter into an arbitrarily risky position for which the estimated value-at-risk is zero.

Clearly it is unreasonable to think that a trader or trading desk might have access to all markets, so the case of a trader who is able to enter into a risky position with an estimated value-at-risk of zero because the estimated covariance matrix he faces is singular is not realistic.[2] However, a U.S. dollar-based interest rate swaps trader will be able to execute transactions in all segments of the U.S. dollar yield curve, and in many cases a corporate end-user's risk management staff will be able to execute transactions at essentially any maturity in several of the actively traded currencies. In the context of risk management systems, these situations correspond to K equal to approximately 20,[3] and K between 50 and 100, or even greater, respectively. It turns out that even in these realistic situations the estimated covariance matrix can sometimes be close to non-singular, and the estimated value at risk can be a very badly biased estimate of the "true" value at risk. Specifically, in these situations the expected value of the ratio of the estimated to "true" value-at-risk, that is the expected value of the ratio

$$\frac{\text{estimated value-at-risk}}{\text{"true" value-at-risk}} = \frac{k\sqrt{w'\hat{\Sigma}_N w}}{k\sqrt{w'\Sigma w}},$$

can be much smaller than one.

[2] However, one of the authors is aware of an international bank whose proprietary trading group is permitted to trade in essentially all markets (subject to position limits).

[3] In some value-at-risk systems, the yield curve for each currency is summarized in terms of approximately 20 basic or "standardized" positions, and an actual instrument or portfolio (e.g., an interest rate swap or a trading "book" of swaps) is interpreted as a portfolio of the 20 standardized positions. From the perspective of the risk measurement system, a fixed income trader is just working with portfolios of these 20 standardized positions.

II. Maximum Bias in Estimated Value-at-Risk

It is frequently suggested that value-at-risk can be used for monitoring and controlling traders and trading desks. For example, risk or position limits might be expressed in terms of value-at-risk, with the value-at-risk then monitored daily (or perhaps more frequently) for violations of the limits. In this "control" context, it is interesting to see whether estimation errors in value-at-risk due to sampling variation in the estimated covariance matrix allow the trader or trading desk to exceed the risk limits.

To address this question, we consider a trader who seeks to evade risk limits and take on as much risk as possible. This may be due to hubris, a desire to exploit convexities in the compensation formula and take advantage of the "trader's option," or simply because she is gambling desperately in an attempt to recover previous losses. We assume that the trader knows the true covariance matrix Σ. It is reasonable to assume that she has a better estimate than $\hat{\Sigma}$, because she is likely to know whether the period from which the *N* observations used to estimate Σ were taken is typical (i.e., she has prior beliefs and access to other information). Among other information, she is likely to have access to market implied volatilities, and perhaps some information from which she can imply market estimates of certain correlations. Assuming knowledge of Σ is the extreme case of assuming that the trader knows more than $\hat{\Sigma}$, and allows us to determine the maximum bias in estimated value-at-risk.

Maximizing true value-at-risk subject to a constraint on estimated value-at-risk yields the same portfolio as minimizing estimated value-at-risk subject to a constraint on true value-at-risk. Using this fact, we consider the problem

$$\min_{w} \sqrt{w'\hat{\Sigma}w} \qquad \text{subject to} \qquad \sqrt{w'\Sigma w} = c\,.$$

Letting w^* denote the solution, from the first order conditions it follows immediately that the estimated value-at-risk is $\sqrt{w^{*\prime}\hat{\Sigma}w^*}$, and the ratio of estimated to "true" value-at-risk is

$$R_1(\hat{\Sigma}) = \frac{k\sqrt{w^{*\prime}\hat{\Sigma}w^*}}{k\sqrt{w^{*\prime}\Sigma w^*}} = \frac{1}{c}\sqrt{w^{*\prime}\hat{\Sigma}w^*}\,.$$

Without loss of generality, we let c=1, so this becomes:

$$R_1(\hat{\Sigma}) = \sqrt{w^{*\prime}\hat{\Sigma}w^{*}} \ .$$

We seek the distribution of the ratio $R_1(\hat{\Sigma})$.[4]

In the appendix we show that $R_1(\hat{\Sigma})$ is the square root of the minimal eigenvalue of $\hat{I}$, where $\hat{I}$ is the estimated covariance matrix constructed using a sample of N vectors z_n drawn from a multivariate Normal distribution with a mean of zero and covariance matrix I. That is, letting $\Sigma^{1/2}$ be the symmetric square root of Σ, the vector z_n defined $z_n \equiv \Sigma^{-1/2}x_n$ is distributed multivariate Normal with a covariance matrix I, and $\hat{I} = \sum_{n=1}^{N} l_n z_n z_n'$. The result that $R_1(\hat{\Sigma})$ is the square root of the minimal eigenvalue of $\hat{I}$ means that the distribution of the ratio of the estimated to "true" value-at-risk doesn't depend on Σ, but only on K and N. An immediate implication of this is that the bias in the estimated value-at-risk does not depend on the covariance matrix Σ.

Computation of eigenvalues is straightforward and relatively fast, so the characterization of $R_1(\hat{\Sigma})$ as the square root of the minimal eigenvalue of $\hat{I}$ allows us to simulate its distribution fairly easily. We draw a sample of N random vectors z_n, construct the estimated covariance matrix $\hat{I} = \sum_{n=1}^{N} l_n z_n z_n'$, and compute the minimal eigenvalue of $\hat{I}$ and its square root. Repeating this process allows us to simulate the distribution of $R_1(\hat{\Sigma})$.

Table 1 shows the mean and standard deviation of $R_1(\hat{\Sigma})$ for K = 10, 20, 50, and 100 when $\hat{I}$ is estimated using the equally weighted covariance matrix estimator

$$\hat{I} = \tfrac{1}{N}\sum_{n=1}^{N} z_n z_n'$$

and N = 50, 100, 200, 500, and 1000. Each case is estimated using 1000 simulated realizations of $R_1(\hat{\Sigma})$. In addition, the first two rows of the table show the mean and standard deviation of

[4] Below we report statistics of the distribution of the ratio of estimated to true value-at-risk rather than the ratio of true to estimated value-at-risk because the estimated value-at-risk is often close to zero, distorting some of the statistics.

R_1 when the covariance matrix is estimated using the exponentially weighted covariance matrix estimator

$$\hat{I} = (1- l)\sum_{n=1}^{N} l^{n-1} z_n z_n' ,$$

where the sum is truncated at N = 100. This exponential weighting scheme has the effect of giving more weight to recent price changes, and is the approach used in J.P. Morgan's RiskMetrics™ methodology. Like them, we set $l = .94$. The results in the table show that the bias is large except when the number of available instruments is small and the number of observations used in estimating the covariance matrix is large. For example, when K = 50 and N = 200, the average ratio of estimated to true value at risk is 0.518. Even when N = 1000, which corresponds to using about 4 years of daily data to estimate the covariance matrix, when K = 50 the average ratio of estimated to true value at risk is 0.786. Moreover, the bias is very large for the exponentially weighted covariance matrix estimator. Even when K is only 10 the mean ratio of estimated to true value at risk is 0.551, and when K = 100 it is only 0.029, that is estimated value-at-risk is typically only 2.9 percent of true value-at-risk.

To interpret these values of K, note that in value-at-risk systems, it is common to summarize the yield curve in each currency in terms of approximately 20 basic or "standardized" positions, and an actual instrument (e.g., an interest rate swap) is interpreted as a portfolio of the 20 standardized positions (e.g., Guldimann 1994). From the perspective of the risk measurement system, a fixed income trader is just working with portfolios of these 20 standardized positions. Thus, K = 20 corresponds to a trader or trading desk which trades the entire yield curve in one currency, e.g. a swaps trading desk, while K = 50 and K = 100 correspond to trading the yield curves in 2 to 3 and 5 to 6 currencies, respectively. These latter cases correspond to the treasury of a corporate end-user of derivatives which actively manages positions in several currencies.

The standard deviations of R_1 reported in Table 1 indicate that the ratios are relatively tightly clustered about the mean values reported in the table. This is confirmed by Table 2, which reports various percentiles of the distributions of R_1, and also the maximums and minimums. The medians in this table are close to the means in Table 1, indicating that the means provide a reasonable measure of the center of the distributions of the ratios of estimated to true value-at-risk. Strikingly, even many of the maximum values are relatively small. For example, when K = 50 and N = 200, the maximum ratio (of 1000) of estimated to true value at risk is 0.569, only slightly higher than the mean of 0.518. Even when N = 1000, when K = 50 the maximum ratio of estimated to true value-at-risk is 0.812. Also, as one might expect after examining the means, the maximums are strikingly small for the exponentially weighted covariance matrix estimator. Even when K is only 10 the maximum ratio of estimated to true value at risk is 0.696, and when $K = 100$ it is only 0.036.

These results raise concerns about the ability of risk limits based on value-at-risk to control the risk-taking behavior of a trader who consciously seeks to evade them. If value-at-risk is to be used for this purpose, the results in Tables 1 and 2 suggest that the covariance matrix should be estimated using a large sample of past price changes. Alternatively, because the bias does not depend on the covariance matrix Σ, the results in this table allow one to adjust conventional value-at-risk estimates to compute estimates based on the assumption that the trader seeks to evade risk limits and maximize the risk of the position.

We emphasize that these measures of bias represent (simulation estimates of) upper bounds on the bias in estimated value-at-risk. In these calculations, we assume that the trader seeks to evade risk limits and take on as much risk as possible, and assume that the trader knows the true covariance matrix Σ. If the trader had a better estimate than $\hat{\Sigma}$, but did not know Σ, these upper bounds would not be reached. Also, our analysis does not consider other mechanisms to control risk-taking such as position limits on individual instruments. Nonetheless, in considering the use of value-at-risk in the "control" function, it is reasonable to consider the worst case. These results call into question the use of value-at-risk for controlling the risk-taking behavior of individual traders or trading desks.

III. Trader maximizes expected return subject to a constraint on estimated value-at-risk

A situation that arises naturally is that of a trader who maximizes expected return subject to the constraint that estimated value-at-risk may not exceed some maximum limit. In this case it is likely that the "true" value-at-risk will exceed the limit, because the trader will tend to take risky positions for which the estimated value-at-risk is underestimated. This can happen even if the trader makes no effort to exceed the limit on value-at-risk, because the goal of maximizing expected return rewards the trader for taking on risk, but the constraint penalizes her for taking on *estimated* risk. Of course, in this case the bias will be less than the maximum bias above.

Letting m denote a $K \times 1$ vector of expected returns, the problem of maximizing expected return subject to a constraint on estimated value-at-risk is

$$\max_{w} w'm \quad \text{subject to} \quad \sqrt{w'\hat{\Sigma}w} \le c\,.$$

The first order conditions are

$$\begin{aligned} m - 2\, l\hat{\Sigma}w &= 0, \\ w'\Sigma w - c &= 0. \end{aligned}$$

The portfolio the trader will choose is $w^* = \dfrac{\sqrt{c}\hat{\Sigma}^{-1}m}{\sqrt{m'\hat{\Sigma}^{-1}m}}$, and the estimated value-at-risk of her portfolio is $k\sqrt{c}$. However, the "true" value-at-risk will be

$$k\sqrt{w^{*\prime}\Sigma w^*} = \frac{k\sqrt{c}\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}{\sqrt{m'\hat{\Sigma}^{-1}m}}\,.$$

The ratio of estimated to "true" value-at-risk is

$$\begin{aligned} R_2(\hat{\Sigma}) &= k\sqrt{c} \Bigg/ \frac{k\sqrt{c}\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}{\sqrt{m'\hat{\Sigma}^{-1}m}} \\ &= \frac{\sqrt{m'\hat{\Sigma}^{-1}m}}{\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}\,. \end{aligned}$$

In the appendix we show that the distribution of R_2 doesn't depend on Σ, but rather depends only on *K* and *N*. Thus, we can simulate it using draws from a multivariate Normal distribution with a covariance matrix of *I*.

Table 3 shows the mean and standard deviation of $R_2(\hat{\Sigma})$ for the same choices of *K* and *N* used in Tables 1 and 2, along with results when the covariance matrix is estimated using the exponentially weighted estimator. As in Tables 1 and 2, each case is estimated using 1000 simulated realizations of $R_2(\hat{\Sigma})$. The biases are considerably smaller than in Table 1. For example, when $K = 50$ and $N = 200$, the average ratio of estimated to true value at risk is 0.753, in comparison to the average ratio of 0.518 in Table 1. When $N = 1000$ and $K = 50$ the average ratio of estimated to true value at risk is 0.951 rather than the 0.786 in Table 1. However, the bias is still large for some combinations of K and N, and is very large for the exponentially weighted covariance matrix estimator except when $K = 10$. With the exponentially weighted estimator the mean ratio of estimated to true value at risk is 0.206 when $K = 50$, and 0.046 when $K = 100$.

Table 4 reports various percentiles of the distributions of $R_2(\hat{\Sigma})$, and also the maximums and minimums. As was the case with Tables 1 and 2, the medians in this table are close to the means reported in Table 3. Again as one might expect after examining the means, even some of the maximums are small for the exponentially weighted covariance matrix estimator. When $K = 20$ the maximum ratio of estimated to true value at risk is 0.793, and when $K = 100$ it is only 0.083.

The biases reported in Table 3 are considerably smaller than those reported in Table 1. This is not surprising, given that Table 1 reports upper bounds on the biases, and assumes that the trader maximizes risk. In contrast, Table 3 reports biases that may well be typical. They stem from maximizing expected return, subject to a constraint on estimated value-at-risk, and do not assume that the trader has knowledge of the true covariance matrix. This may be a reasonable approximation of the behavior of traders who face a constraint on estimated value-at-risk; finance theory suggests that it is. Thus, the results reported in Tables 3 and 4 raise doubts about the use of value-at-risk in monitoring and performance evaluation when $K \geq 50$ and N is small or moderate in size, or when the exponentially weighted covariance estimator is used, regardless of the value of K.

IV. Trader chooses a portfolio as close as possible to a desired portfolio subject to a constraint on estimated value-at-risk

Traders may not always maximize expected return. Rather, there may be some idiosyncratic trade that, for whatever reason, the trader or trading desk wants to do. This is the case we consider in this section.

Specifically, we consider a portfolio $\overline{w}$, which will be interpreted as a portfolio that the trader would like to establish. We choose $\overline{w}$ so that the "true" value-at-risk of $\overline{w}$ is greater than the limit imposed by senior management. The portfolio that the trader enters into is denoted w. Senior management imposes the constraint $w'\hat{\Sigma}_N w \leq c$ (the constraint is in terms of the estimated value-at-risk because that is all senior management can observe). The trader tries to get close to $\overline{w}$ without violating the constraint $w'\hat{\Sigma}_N w \leq c$. Formally, she solves the problem

$$\min_{w} \frac{1}{2}(w-\overline{w})'Q(w-\overline{w}) \qquad \text{subject to} \qquad \sqrt{w'\hat{\Sigma}_N w} \le c,$$

where Q is a matrix that weights deviations between w and $\overline{w}$.

A reasonable choice of Q is $Q = \Sigma$, because this choice of Q corresponds to minimizing the variance of the difference between the returns of the portfolios w and $\overline{w}$. To see this, note that the trading desk wants to enter into the portfolio $\overline{w}$, and thus desires the random variable $\overline{w}'x$ (recall that x is the vector of price changes), but is forced to accept a random variable $w'x$. Writing $w'x$ in terms of $\overline{w}'x$, we have $w'x = \overline{w}'x + (w-\overline{w})'x$, where $(w-\overline{w})'x$ is the difference between the returns of the two portfolios. The variance of $(w-\overline{w})'x$ is $(w-\overline{w})'\Sigma(w-\overline{w})$, so the choice of $Q = \Sigma$ corresponds to minimizing this variance. Also, at the optimal w we have $E(w'x|\overline{w}'x) = 0.$

With this choice of Q, the problem becomes

$$\min_{w} \tfrac{1}{2}(w-\overline{w})'\Sigma(w-\overline{w}) \text{subject to} \sqrt{w'\hat{\Sigma}w} \le c.$$

Letting w^* denote the solution, once again we are interested in the ratio of estimated to "true" value-at-risk,

$$R_3(\hat{\Sigma}) = \frac{\sqrt{w^{*\prime}\hat{\Sigma}w^*}}{\sqrt{w^{*\prime}\Sigma w^*}}.$$

Since $\sqrt{w^{*\prime}\hat{\Sigma}w^*} = c$, this becomes $R_3(\hat{\Sigma}) = \dfrac{c}{\sqrt{w^{*\prime}\Sigma w^*}}.$

The interpretation of this case is that the trading desk wants to enter into a portfolio $\overline{w}$, which is too risky in the sense that $\overline{w}'\hat{\Sigma}\overline{w} > c$. So, the trading desk instead enters into the portfolio w^*, which is as close as possible to $\overline{w}$ without violating the constraint. Then, $R_3(\hat{\Sigma})$ indicates the relation between the estimated value-at-risk $k\sqrt{w^{*\prime}\hat{\Sigma}w^*}$ and the "true" value-at-risk $k\sqrt{w^{*\prime}\Sigma w^*}$ and.

We show in the appendix that the distribution of $R_3(\hat{\Sigma})$ depends only on the true value-at-risk $\sqrt{\overline{w}'\Sigma\overline{w}}$ of the desired portfolio $\overline{w}$. This allows use to carry out simulations by letting $\Sigma = I$. Table 5 shows the mean and standard deviation of $R_3(\hat{\Sigma})$ when $c = \sqrt{w'\hat{\Sigma}w} = 1$ and

$\sqrt{w'\Sigma w} = \sqrt{2}$, i.e. the variance of the desired portfolio is twice the permitted variance. The table reports results for the equally weighted covariance matrix estimator for the same choices of *K* and *N* used in Tables 1-4, along with results when the covariance matrix is estimated using the exponentially weighted estimator. As in the other tables, each case is estimated using 1000 simulated realizations of $R_3(\hat{\Sigma})$.

The biases in Table 5 are considerably smaller than those reported in Tables 1 and 3, though still significant for many of the combinations of *K* and *N*. For example, when *K* = 50 and *N* = 200, the average ratio of estimated to true value at risk is 0.889, in comparison to the average ratio of 0.518 in Table 1 and average ratio of 0.753 in Table 3. However, the biases are large only for the exponentially weighted covariance matrix estimator. For example, for the exponentially weighted estimator with *K* = 50, the average ratio is 0.655. This raises concerns about the use of this estimator.

V. Conclusion

We have shown that there can be significant biases in value-at-risk estimates. How should one interpret these results?

First, our results have no implications for many uses of value-at-risk. We study situations in which an individual trader or decision making unit (e.g., a trading desk) either intentionally or unintentionally systematically exploits the estimation errors in value-at-risk in order to enter into positions for which the "true" value-at-risk exceeds the estimated value-at-risk. The systematic exploitation of the estimation errors is crucial; without it, and setting aside other estimation issues not addressed in this paper, the estimated value-at-risk would be an unbiased estimate of "true" value-at-risk. For this reason, our results have implications only for the use of value-at-risk to monitor or control individual traders or trading desks. For firms or companies with multiple trading desks, whether and how the biases at the level of the individual decision making unit aggregate to biases at the firm level will depend upon the correlations among the portfolios chosen by the individual units. Thus, our results do not apply to the use of value-at-risk in reporting a summary measure of aggregate market risk to senior management or the board of directors, investors, or regulators.

However, our results have strong implications for the use of value-at-risk in controlling individual traders or trading desks. We find that the bias in estimated value-at-risk can be large when a trader or trading desk deliberately seeks to evade risk limits and take on as much risk as possible. This raises questions about the efficacy of value-at-risk in controlling the behavior of individual risk-taking units.

In our analysis and simulations, we assume that the covariance matrix is the same every period. Under this assumption, the biases can be mitigated by the use of large samples to estimate value-at-risk. However, in actuality the covariance matrices of price changes are not constant (see Figlewski (1997) and section 12.2 of Cambell, Lo, and MacKinlay (1997), and the references cited therein), and value-at-risk measures that assume that they are can lead to large errors. In addition, work by Kupiec (1995), Lopez (1997), and Crnkovic and Drachman (1996) suggests that some tests for identifying errors in value-at-risk models have relatively little power against reasonable alternatives. Thus, one cannot be confident that systematic biases in value-at-risk estimates will be readily detected.

Our results also have implications for the use of value-at-risk in performance evaluation and compensation. Recently, it has been suggested that compensation should be based on risk-adjusted performance (e.g., Davies (1997)). If the risk adjustment is done using value-at-risk, then traders will have clear incentives to enter into portfolios in which the estimated value-at-risk is low.

References

Alexander, Carol O., "Estimating and Forecasting Volatility and Correlation: Methods and Applications," *Financial Derivatives and Risk Management*, No. 7 (September 1996).

Alexander, Carol O. and C.T. Leigh, "On the Covariance Matrices Used in Value at Risk Models," *Journal of Derivatives* 4, No. 3 (1997), 50-62.

Beder, Tanya S., "VAR: Seductive but Dangerous." *Financial Analysts Journal*, September-October 1995, pp. 12-24.

Boudoukh, J., M. Richardson, and R.F. Whitelaw, "Investigation of a Class of Volatility Estimators," *Journal of Derivatives* 4, No. 3 (1997), pp. 62-89.

Butler, J.S., and B. Schachter, 1996, "Improving Value-at-Risk Estimates by Combining Kernel Estimation with Historical Simulation," working paper, Vanderbilt University.

Campbell, Hohn Y., Andrew M. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey.

Crnkovic, C. and J. Drachman. "Quality Control," *Risk* 9, No. 9 (September 1996), pp. 138-143.

Danielson, Jon and Casper G. de Vries. "Extreme Returns, Tail Estimation, and Value-at-Risk." Working paper, Institute of Economic Studies, University of Iceland, 1997.

Davies, Daniel. "Remuneration and Risk." *Financial Stability Review* 2 (spring 1997), pp. 18-23.

Duffie, D. and J. Pan, "An Overview of Value at Risk," *Journal of Derivatives* 4, No. 3 (1997), pp. 7-49.

Figlewski, Stephen. "Forecasting Volatility." *Financial Markets, Institutions, and Instruments* 6, No. 1. (January 1997), pp. 1-88.

Guldimann, Til. *RiskMetricsÔ Technical Document* (2nd edition). New York: J.P. Morgan 1994.

Hendricks, Darryll, "Evaluation of Value-at-Risk Models Using Historical Data." *Federal Reserve Bank of New York Economic Policy Review* 2, No. 1 (April 1996), pp. 39-62.

Jordan, James V., and Robert J. McKay, "Assessing Value At Risk for Equity Portfolios: Implementing Alternative Techniques," Center for Study of Futures and Options Markets, Virginia Polytechnic Institute, Working paper, July 1995.

Jorion, P., 1996, "Risk^2: Measuring the Risk in Value at Risk," forthcoming, *Financial Analysts Journal*.

Kupiec, Paul. "Techniques for Verifying the Accuracy of Risk Measurement Models." *Journal of Derivatives* 3, No. 2 (Winter 1995), pp. 73-84.

Linsmeier, Thomas J. and Neil D. Pearson, "Risk Measurement: An Introduction to Value at Risk", working paper, Office for Futures and Options Research, University of Illinois, 1996.

Linsmeier, Thomas J. and Neil D. Pearson, "Risk Measurement," in Gary Klopfenstein, ed., *FX: Managing Global Currency Risk*, Glenlake Publishing Co., 1997a.

Linsmeier, Thomas J., and Neil D. Pearson, "Quantitative Disclosures of Market Risk in the SEC Release," *Accounting Horizons* 11, No. 1 (1997b), 107-135.

Longerstaey, Jacques. *RiskMetrics™ Technical Document* (4th edition). New York: j.P. Morgan 1996.

Lopez, Jose A. "Regulatory Evaluation of Value-at-Risk Models." Working paper, Federal Reserve Bank of New York, 1997.

Mahoney, James P. "Forecast Biases in Value-at-Risk Estimations: Evidence from Foreign Exchange and Global Equity Portfolios."

Marshall, Chris, and Michael Siegel, "Value at Risk: Implementing a Risk Measurement Standard." *Journal of Derivatives* 4, No. 3 (1997), pp. 91-111.

Pritzker, Mathew, "Evaluating Value at Risk Methodologies: Accuracy versus Computational Time," working paper, Board of Governors of the Federal Reserve System, 1996.

Robinson, Gary, "More Haste, Less Precision," *Risk* 9, No. 9 (September 1996), pp. 117-121.

Zangari, Peter, "A VaR Methodology for Portfolios that Include Options," *RiskMetricsÔ Monitor*, First Quarter 1996, pp. 4-12.

Appendix

Proof that $R_1(\hat{\Sigma})$ is the square root of the minimal eigenvalue of $\hat{I}$.

Consider the problem

$$\min_{w} w'\hat{\Sigma}w \quad \text{subject to} \quad w'\Sigma w = 1 .$$

Letting $w*$ denote the solution of this problem, the ratio of estimated to “true” value-at-risk is

$$R_1(\hat{\Sigma}) = \frac{k\sqrt{w*'\hat{\Sigma}w*}}{k\sqrt{w*'\Sigma w*}} = \sqrt{w*'\hat{\Sigma}w*}.$$

Forming the Lagrangian and computing the first order conditions, the solution $w*$ must satisfy:

$$(\hat{\Sigma}_N - j\,\Sigma)w* = 0, \tag{5}$$
$$w*'\Sigma w* - 1 = 0, \tag{6}$$

where j is the Lagrange multiplier. Pre-multiplying (5) by $w*'$ and using (6), we obtain

$w*'\hat{\Sigma}_N w* = j$, or

$$R_1(\hat{\Sigma}_N) = \sqrt{w*'\hat{\Sigma}_N w*} = \sqrt{j} \ .$$

This shows that the ratio of the estimated to true value-at-risk is the square root of the Lagrange multiplier. In general, there may be multiple solutions of the first order conditions (5) and (6), so we want the solution with the smallest j .

Each j that solves (5) and (6) is an eigenvalue of $\hat{I}$. To see this, pre-multiply (5) by Σ^{-1} , yielding

$$(\Sigma^{-1}\hat{\Sigma} - j\,)w* = 0. \tag{7}$$

The matrix $\hat{\Sigma}$ is estimated by $\hat{\Sigma} = \sum_{n=1}^{N} l_n x_n x_n'$, where $x_n = \Sigma^{1/2} z_n$, $z_n \sim N(0, I_K)$, and I_K is the K-dimensional identity matrix. Noticing that $\hat{I} = \sum_{n=1}^{N} l_n z_n z_n'$, we have

$$\hat{\Sigma} = \Sigma^{1/2}\hat{I}\,\Sigma^{1/2} \ . \tag{8}$$

Substituting (8) into (7) yields $(\Sigma^{-1/2}\hat{I}\,\Sigma^{-1/2} - j\,)w* = 0$. This is the standard eigenvalue equation, so each j that solves this equation is an eigenvalue of $\Sigma^{-1}\hat{I}\Sigma^{-1}$, and $w*$ is the associated eigenvector. Now we need only show that the two matrices $\Sigma^{-1/2}\hat{I}\,\Sigma^{-1/2}$ and $\hat{I}$ share the same eigenvalues.

Suppose j is an eigenvalue of $\Sigma^{-1/2}\hat{I}\,\Sigma^{-1/2}$ and w is the associated eigenvector, i.e.

$$\Sigma^{-1/2}\hat{I}\,\Sigma^{-1/2}w = j\,w\,.$$

Pre-multiplying this equation by $\Sigma^{1/2}$ yields $\hat{I}\,\Sigma^{-1/2}w = j\,\Sigma^{1/2}w$. Then defining $v = \Sigma^{1/2}w$, we have $\hat{I}v = j\,v$. Thus j is also an eigenvalue of $\hat{I}$. This argument can be reversed to show that if j is an eigenvalue of $\hat{I}$, then it is also an eigenvalue of $\Sigma^{-1/2}\hat{I}\,\Sigma^{-1/2}$.

Proof that the distribution of $R_2(\hat{\Sigma})$ depends only on *K* and *N*.

The problem of maximizing expected return subject to a constraint on the estimated value-at-risk is

$$\max_w w'm \qquad \text{subject to} \qquad w'\hat{\Sigma}\,w \le c\,.$$

The optimal portfolio w^* satisfies the first order conditions

$$m - 2\,l\hat{\Sigma}w^* = 0,$$
$$w^{*\prime}\Sigma w^* - c = 0.$$

Rearranging the first equation, we have

$$w^* = \frac{\hat{\Sigma}^{-1}m}{2\,l}\,. \tag{9}$$

Substituting this into the second of the first order conditions and rearranging, we obtain $\frac{1}{2\,l} = \frac{\sqrt{c}}{\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}$. Then using the fact that $\hat{\Sigma} = \Sigma^{1/2}\hat{I}\Sigma^{1/2}$, we have $w^* = \frac{\sqrt{c}\hat{\Sigma}^{-1}m}{\sqrt{m'\hat{\Sigma}_N^{-1}m}}$.

The estimated value-at-risk of this portfolio is $k\sqrt{w^{*\prime}\hat{\Sigma}w^*} = k\sqrt{c}$, and the "true" value-at-risk is $k\sqrt{w^{*\prime}\Sigma w^*} = \frac{k\sqrt{c}\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}{\sqrt{m'\hat{\Sigma}^{-1}m}}$. Thus the ratio of estimated to true value-at-risk is

$$R_2(\hat{\Sigma}) = \frac{\sqrt{m'\hat{\Sigma}^{-1}m}}{\sqrt{m'\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}m}}.$$

The distribution of this ratio depends only on *K* and *N*. To see this, from equation (8) and $\hat{I} = \sum_{n=1}^{N} l_n z_n z_n'$, we have $\hat{\Sigma}^{-1} = \Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}$, so the ratio of the estimated to true value-at-risk becomes

$$R_2(\hat{\Sigma}) = \frac{\sqrt{m'\Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}m}}{\sqrt{m'\Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}\Sigma\Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}m}} = \sqrt{\frac{m'\Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}m}{m'\Sigma^{-1/2}\hat{I}^{-1}\hat{I}^{-1}\Sigma^{-1/2}m}}$$

Let *T* be the orthogonal matrix such that $\Sigma^{-1/2}m = aTv,$ where $v = (1,0,\dots,0)'$ and $a^2 = m'\Sigma^{-1}m$. Then

$$\begin{aligned}
\frac{m'\Sigma^{-1/2}\hat{I}^{-1}\Sigma^{-1/2}m}{m'\Sigma^{-1/2}\hat{I}^{-1}\hat{I}^{-1}\Sigma^{-1/2}m} &= \frac{v'T'\hat{I}^{-1}Tv}{v'T'\hat{I}^{-1}\hat{I}^{-1}Tv} \\
&= \frac{v'T'\hat{I}^{-1}Tv}{v'T'\hat{I}^{-1}TT'\hat{I}^{-1}Tv} \\
&= \frac{v'(T'\hat{I}T)^{-1}v}{v'(T'\hat{I}T)^{-1}(T'\hat{I}T)^{-1}v} \\
&= \frac{v'(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}v}{v'(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}v}.
\end{aligned}$$

Since the distribution of $T'z_n$ is identical to the distribution of z_n, the distribution of

$$\frac{v'(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}v}{v'(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}(T'(\sum_{n=1}^{N} l_n z_n z_n')T)^{-1}v}$$

is identical to the distribution of

$$\frac{v'(\sum_{n=1}^{N} l_n z_n z_n')^{-1}v}{v'(\sum_{n=1}^{N} l_n z_n z_n')^{-1}(\sum_{n=1}^{N} l_n z_n z_n')^{-1}v}$$

, or more simply, $\frac{v'\hat{I}^{-1}v}{v'\hat{I}^{-1}\hat{I}^{-1}v}$, which only varies with the choice of *N* and *K*. Thus we conclude that the distribution of the ratio $R_2(\hat{\Sigma})$ depends only on *K* and *N*.

Proof that the distribution of $R_3(\hat{\Sigma})$ depends only on the "true" value-at-risk $\sqrt{\overline{w}'\Sigma\overline{w}}$.

The problem is:

$$\min_w \frac{1}{2}(w-\overline{w})'\Sigma(w-\overline{w}) \qquad \text{subject to} \qquad w'\hat{\Sigma}w \le c. \tag{10}$$

Letting w^* denote the solution of (5) and $R_3(\hat{\Sigma})$ denote the ratio of estimated to true value-at-risk, we have

$$R_3(\hat{\Sigma}) = \frac{\sqrt{w^{*\prime}\hat{\Sigma}w^*}}{\sqrt{w^{*\prime}\Sigma w^*}} = c\Big/\sqrt{w^{*\prime}\Sigma w^*}.$$

Equation (8) above tells us that $\hat{\Sigma}$ is a function of Σ and the z_n, so w^* and $R_3(\hat{\Sigma})$ are functions of Σ, $\overline{w}$, and the z_n. We use $w^*(\Sigma,\overline{w},z_1,\ldots,z_N)$ and $R_3(\Sigma,\overline{w},z_1,\ldots,z_N)$ to denote them.

Suppose that there are two portfolios $\overline{w}_1$ and $\overline{w}_2$ with the same true value-at-risk, i.e., $\overline{w}_1'\Sigma\overline{w}_1 = \overline{w}_2'\Sigma\overline{w}_2$ or $\left\|\Sigma^{1/2}\overline{w}_1\right\|^2 = \left\|\Sigma^{1/2}\overline{w}_2\right\|^2$. This implies that there exists an orthogonal matrix T such that

$$T\Sigma^{1/2}\overline{w}_2 = \Sigma^{1/2}\overline{w}_1. \tag{11}$$

For $\overline{w}_1$, using equation (8), the problem (12) can be written

$$\min_w \left\|\Sigma^{1/2}w - \Sigma^{1/2}\overline{w}_1\right\|^2 \qquad \text{subject to} \qquad w'\Sigma^{1/2}\Big(\sum_{n=1}^{N} l_n z_n z_n'\Big)\Sigma^{1/2}w = c. \tag{12}$$

If $w^*(\Sigma,\overline{w}_1,z_1,\ldots,z_N)$ is the solution of this problem, then by equation (11) it is also the solution of

$$\min_w \left\|\Sigma^{1/2}w - T\Sigma^{1/2}\overline{w}_2\right\|^2 \qquad \text{subject to} \qquad w'\Sigma^{1/2}\Big(\sum_{n=1}^{N} l_n z_n z_n'\Big)\Sigma^{1/2}w = c.$$

Since T' is orthogonal, $\left\|\Sigma^{1/2}w - T\Sigma^{1/2}\overline{w}_2\right\|^2 = \left\|T'(\Sigma^{1/2}w - T\Sigma^{1/2}\overline{w}_2)\right\|^2$. Then the problem becomes

$$\min_w \left\|T'\Sigma^{1/2}w - \Sigma^{1/2}\overline{w}_2\right\|^2 \qquad \text{subject to} \qquad w'\Sigma^{1/2}\Big(\sum_{n=1}^{N} l_n z_n z_n'\Big)\Sigma^{1/2}w = c.$$

Letting v be the vector $v = \Sigma^{-1/2}T'\Sigma^{1/2}w$, we have $T'\Sigma^{1/2}w = \Sigma^{1/2}v$, and this can be rewritten as

$$\min_v \left\| \Sigma^{1/2} v - \Sigma^{1/2} \overline{w}_2 \right\|^2 \quad \text{subject to} \quad v' \Sigma^{1/2} T' (\sum_{n=1}^{N} l_n z_n z_n') T \Sigma^{1/2} v = c .$$

Letting v^* denote the solution of this problem, by comparing this problem with (12), we know $v^* = w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N)$. Therefore,

$$w^*(\Sigma, \overline{w}_1, Z) = \Sigma^{-1/2} T \Sigma^{1/2} v^* = \Sigma^{-1/2} T \Sigma^{1/2} w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N),$$

where the first equality follows from the definition of *v*. Using this equation, we have

$$\begin{aligned} R_3(\Sigma, \overline{w}_1, z_1, \ldots, z_N) &= 1 \Big/ \sqrt{w^*(\Sigma, \overline{w}_1, z_1, \ldots, z_N)' \Sigma w^*(\Sigma, \overline{w}_1, z_1, \ldots, z_N)} \\ &= 1 \Big/ \sqrt{w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N)' \Sigma^{1/2} T' \Sigma^{-1/2} \Sigma \Sigma^{-1/2} T \Sigma^{1/2} w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N)} \\ &= 1 \Big/ \sqrt{w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N)' \Sigma w^*(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N)} \\ &= R_3(\Sigma, \overline{w}_2, T'z_1, \ldots, T'z_N). \end{aligned}$$

Since z_n and $T'z_n$ are two equally likely realizations of $N(0, I_K)$, the distribution of $R_3(\hat{\Sigma})$ for $\overline{w}_1$ is identical to that for $\overline{w}_2$. Then, since $\overline{w}_1$ and $\overline{w}_2$ are arbitrarily selected with the only requirement being that they have the same true value-at-risk, the distribution of $R_3(\hat{\Sigma})$ depends on $\overline{w}$ only through $\sqrt{\overline{w}' \Sigma \overline{w}}$. Next, we will show that it depends on Σ only through $\sqrt{\overline{w}' \Sigma \overline{w}}$.

Suppose $w^*(\Sigma, \overline{w}, z_1, \ldots, z_N)$ solves the problem

$$\min_w \left\| \Sigma^{1/2} w - \Sigma^{1/2} \overline{w} \right\|^2 \quad \text{subject to} \quad w' \Sigma^{1/2} (\sum_{n=1}^{N} l_n z_z z_n') \Sigma^{1/2} w = c .$$

Letting $v = \Sigma^{1/2} w$, the problem can be rewritten as

$$\min_v \left\| v - \Sigma^{1/2} \overline{w} \right\|^2 \quad \text{subject to} \quad v' (\sum_{n=1}^{N} l_n z_n z_n') v = c .$$

If v^* solves this problem then, by comparing the problem with (7), we have $v^* = w^*(I_K, \Sigma^{1/2} \overline{w}, z_1, \ldots, z_N)$. Thus,

$$w^*(\Sigma, \overline{w}, z_1, \ldots, z_N) = \Sigma^{-1/2} v^* = \Sigma^{-1/2} w^*(I_K, \Sigma^{1/2} \overline{w}, z_1, \ldots, z_N),$$

where I_K is the *K*-dimensional identity matrix. Using this equation, we have

$$
\begin{aligned}
R_3(\Sigma,\overline{w},z_1,\ldots,z_N) &= 1\Big/\sqrt{w^*(\Sigma,\overline{w},z_1,\ldots,z_N)'\Sigma w^*(\Sigma,\overline{w},z_1,\ldots,z_N)} \\
&= 1\Big/\sqrt{w^*(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N)'\Sigma^{-1/2}\Sigma\Sigma^{-1/2}w^*(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N)} \\
&= 1\Big/\sqrt{w^*(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N)'I_K w^*(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N)} \\
&= R_3(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N).
\end{aligned}
$$

From the previous argument we know that the distribution of $R_3(I_K,\Sigma^{1/2}\overline{w},z_1,\ldots,z_N)$ is a function of only I_K and the true value-at-risk of the objective portfolio, $\overline{w}'\Sigma\ \overline{w}$. Thus, the distribution of $R_3(\Sigma,\overline{w},z_1,\ldots,z_N)$ is also a function of only I_K and $\overline{w}'\Sigma\ \overline{w}$. Put another way, the distribution of $R_3(\Sigma,\overline{w},Z)$ depends on Σ only through $\overline{w}'\Sigma\ \overline{w}$.

Table 1: Means and standard deviations of $R_1(\hat{\Sigma})$, the ratio of estimated to true value-at-risk assuming that the trader maximizes the bias in the estimated value-at-risk. The standard deviations are in parentheses. The distributions were estimated using 1000 simulated realizations of $R_1(\hat{\Sigma})$.

	Dimension of Covariance Matrix (*K*)			
Number of observations used to estimate covariance matrix (*N*)	10	20	50	100
exponential weighting with λ = .94 and *N* = 100[5]	0.551 (0.047)	0.372 (0.031)	0.131 (0.011)	0.029 (0.002)
50	0.606 (0.053)	0.405 (0.042)		
100	0.725 (0.039)	0.586 (0.032)	0.312 (0.023)	
200	0.809 (0.029)	0.710 (0.025)	0.518 (0.019)	0.306 (0.014)
500	0.879 (0.019)	0.817 (0.016)	0.697 (0.013)	0.563 (0.011)
1000	0.915 (0.013)	0.871 (0.012)	0.786 (0.009)	0.692 (0.008)

[5] When K = 100, N = 200 observations were used.

Table 2: Percentiles of the distributions of $R_1(\hat{\Sigma})$, the ratio of estimated to true value-at-risk assuming that the trader maximizes the bias in the estimated value-at-risk. The distributions were estimated using 1000 simulated realizations of $R_1(\hat{\Sigma})$.

Number of observations used to estimate covariance matrix (*N*)	Percentile	Dimension of Covariance Matrix (*K*) 10	20	50	100
exponential weighting with λ = .94 and *N* = 100[6]	min.	0.406	0.277	0.092	0.022
	10th	0.490	0.333	0.117	0.026
	25th	0.520	0.351	0.124	0.027
	50th	0.554	0.374	0.132	0.029
	75th	0.581	0.394	0.139	0.030
	90th	0.612	0.409	0.145	0.032
	max.	0.696	0.481	0.164	0.036
50	min.	0.424	0.220		
	10th	0.539	0.352		
	25th	0.571	0.378		
	50th	0.607	0.405		
	75th	0.643	0.433		
	90th	0.674	0.456		
	max.	0.760	0.536		
100	min.	0.592	0.485	0.221	
	10th	0.675	0.543	0.282	
	25th	0.700	0.563	0.298	
	50th	0.728	0.588	0.312	
	75th	0.754	0.609	0.328	
	90th	0.774	0.627	0.342	
	max.	0.864	0.676	0.377	
200	min.	0.695	0.611	0.464	0.255
	10th	0.772	0.680	0.494	0.288
	25th	0.791	0.693	0.506	0.297
	50th	0.810	0.712	0.519	0.306
	75th	0.829	0.727	0.531	0.315
	90th	0.846	0.740	0.543	0.323
	max.	0.884	0.781	0.569	0.351
500	min.	0.783	0.767	0.643	0.527
	10th	0.855	0.795	0.681	0.548
	25th	0.866	0.805	0.688	0.556
	50th	0.880	0.818	0.697	0.563
	75th	0.892	0.828	0.706	0.571
	90th	0.903	0.836	0.713	0.577
	max.	0.931	0.861	0.74	0.594
1000	min.	0.869	0.829	0.751	0.666
	10th	0.899	0.856	0.774	0.682
	25th	0.907	0.863	0.781	0.687
	50th	0.915	0.871	0.787	0.693
	75th	0.924	0.879	0.793	0.698
	90th	0.932	0.886	0.798	0.702
	max.	0.966	0.908	0.812	0.714

[6] When *K* = 100, *N* = 200 observations were used.

Table 3: Mean and standard deviation of the distributions of $R_2(\hat{\Sigma})$, the ratio of estimated to "true" value-at-risk assuming that the trader maximizes expected return subject to a constraint on the estimated value-at-risk. The distributions were estimated using 1000 simulated realizations of $R_2(\hat{\Sigma})$. The standard deviations are in parentheses.

	Dimension of Covariance Matrix (*K*)			
Number of observations used to estimate covariance matrix (*N*)	10	20	50	100
exponential weighting with $\lambda = .94$ and $N = 100$[7]	0.742 (0.102)	0.541 (0.075)	0.206 (0.032)	0.046 (0.007)
50	0.814 (0.095)	0.609 (0.092)		
100	0.905 (0.070)	0.805 (0.071)	0.497 (0.061)	
200	0.954 (0.050)	0.901 (0.051)	0.753 (0.048)	0.503 (0.045)
500	0.981 (0.032)	0.959 (0.031)	0.902 (0.033)	0.804 (0.031)
1000	0.991 (0.022)	0.981 (0.022)	0.951 (0.023)	0.901 (0.023)

[7] When $K = 100$, $N = 200$ observations were used.

Table 4: Percentiles of the distributions of $R_2(\hat{\Sigma})$, the ratio of estimated to "true" value-at-risk assuming that the trader maximizes expected return subject to a constraint on the estimated value-at-risk. The distributions were estimated using 1000 simulated realizations of $R_2(\hat{\Sigma})$

Number of observations used to estimate covariance matrix (N)	Percentile	Dimension of Covariance Matrix (K): 10	20	50	100
exponential weighting with λ = .94 and N = 100[8]	min.	0.466	0.348	0.117	0.029
	10^{th}	0.614	0.446	0.166	0.038
	25^{th}	0.671	0.488	0.184	0.041
	50^{th}	0.736	0.534	0.204	0.045
	75^{th}	0.8100	0.589	0.227	0.050
	90^{th}	.876	0.641	0.248	0.054
	max.	1.103	0.793	0.316	0.083
50	min.	0.545	0.376		
	10^{th}	0.698	0.489		
	25^{th}	0.747	0.547		
	50^{th}	0.809	0.608		
	75^{th}	0.877	0.666		
	90^{th}	0.939	0.731		
	max.	1.188	0.911		
100	min.	0.669	0.601	0.339	
	10^{th}	0.819	0.716	0.423	
	25^{th}	0.858	0.756	0.457	
	50^{th}	0.906	0.800	0.494	
	75^{th}	0.951	0.850	0.540	
	90^{th}	0.995	0.899	0.580	
	max.	1.144	1.050	0.702	
200	min.	0.800	0.754	0.597	0.375
	10^{th}	0.891	0.834	0.693	0.445
	25^{th}	0.918	0.869	0.720	0.471
	50^{th}	0.954	0.899	0.753	0.502
	75^{th}	0.986	0.935	0.786	0.535
	90^{th}	1.014	0.964	0.813	0.560
	max.	1.152	1.075	0.939	0.670
500	min.	0.879	0.869	0.789	0.707
	10^{th}	0.938	0.920	0.859	0.764
	25^{th}	0.959	0.937	0.880	0.782
	50^{th}	0.982	0.959	0.902	0.804
	75^{th}	1.001	0.980	0.926	0.825
	90^{th}	1.021	1.000	0.945	0.843
	max.	1.101	1.064	1.007	0.918
1000	min.	0.925	0.914	0.862	0.797
	10^{th}	0.963	0.953	0.922	0.873
	25^{th}	0.976	0.966	0.935	0.886
	50^{th}	0.991	0.982	0.950	0.902
	75^{th}	1.006	0.997	0.965	0.917
	90^{th}	1.021	1.011	0.981	0.929
	max.	1.068	1.054	1.018	0.973

[8] When K = 100, N = 200 observations were used.

Table 5: Mean and standard deviation of the distributions of $R_3(\hat{\Sigma})$, the ratio of estimated to "true" value-at-risk assuming that the trader selects a preferred portfolio subject to a constraint on the estimated value-at-risk. The standard deviations are in parentheses. The distributions were estimated using 1000 simulated realizations of $R_3(\hat{\Sigma})$.

	Dimension of Covariance Matrix (*K*)			
Number of observations used to estimate covariance matrix (*N*)	10	20	50	100
exponential weighting with λ = .94 and *N* = 100[9]	0.868 (0.101)	0.779 (0.085)	0.655 (0.050)	0.583 (0.026)
50	0.912 (0.093)	0.831 (0.084)		
100	0.953 (0.070)	0.908 (0.067)	0.799 (0.055)	
200	0.978 (0.047)	0.955 (0.048)	0.889 (0.046)	0.800 (0.039)
500	0.992 (0.032)	0.982 (0.031)	0.953 (0.030)	0.907 (0.030)
1000	0.997 (0.022)	0.991 (0.023)	0.976 (0.022)	0.952 (0.022)

[9] When *K* = 100, *N* = 200 observations were used.

Table 6: Percentiles of the distributions of $R_3(\hat{\Sigma})$, the ratio of estimated to "true" value-at-risk assuming that the trader selects a preferred portfolio subject to a constraint on the estimated value-at-risk. The distributions were estimated using 1000 simulated realizations of $R_3(\hat{\Sigma})$.

		Dimension of Covariance Matrix (K)			
Number of observations used to estimate covariance matrix (N)	Percentile	10	20	50	100
exponential weighting with λ = .94 and N = 100	min.	0.599	0.348	0.524	0.512
	10th	0.751	0.446	0.596	0.552
	25th	0.791	0.488	0.618	0.564
	50th	0.823	0.534	0.649	0.581
	75th	0.9301	0.589	0.686	0.598
	90th	.001	0.641	0.721	0.616
	max.	1.267	0.793	0.830	0.681
50	min.	0.619	0.593		
	10th	0.792	0.725		
	25th	0.848	0.771		
	50th	0.910	0.826		
	75th	0.974	0.887		
	90th	1.035	0.942		
	max.	1.284	1.100		
100	min.	0.766	0.693	0.657	
	10th	0.860	0.822	0.730	
	25th	0.908	0.863	0.769	
	50th	0.952	0.906	0.795	
	75th	1.002	0.953	0.836	
	90th	1.041	0.995	0.870	
	max.	1.210	1.149	1.009	
200	min.	0.815	0.778	0.749	0.681
	10th	0.919	0.891	0.831	0.750
	25th	0.946	0.923	0.857	0.776
	50th	0.977	0.955	0.888	0.800
	75th	1.007	0.989	0.920	0.824
	90th	1.039	1.016	0.948	0.849
	max.	1.138	1.100	1.028	0.930
500	min.	0.889	0.866	0.865	0.814
	10th	0.952	0.941	0.914	0.869
	25th	0.971	0.961	0.931	0.885
	50th	0.992	0.982	0.952	0.907
	75th	1.013	1.002	0.973	0.927
	90th	1.033	1.023	0.992	0.946
	max.	1.089	1.075	1.048	1.003
1000	min.	0.914	0.924	0.880	0.878
	10th	0.967	0.962	0.949	0.924
	25th	0.982	0.976	0.961	0.938
	50th	0.996	0.989	0.976	0.952
	75th	1.009	1.005	0.991	0.966
	90th	1.025	1.022	1.005	0.980
	max.	1.062	1.055	1.037	1.018