Calculating Quantile-based Risk Analytics with *L*-estimators

Helmut Mausser

Non-parametric methods for calculating quantile-based risk measures (e.g., Value-at-Risk) estimate a quantile of the loss distribution from the order statistics of the sample losses. *L*-estimators, which compute linear combinations of the order statistics, estimate the quantile as a weighted average of one or more losses. Risk management requires estimating not only the quantile, but also its sensitivity to the positions comprising the portfolio. *L*-estimators based on a single order statistic (e.g., the sample quantile) are often inadequate for this purpose; more robust risk analytics can be obtained by computing a weighted average of multiple order statistics. We show that the calculation of marginal VaR and risk contribution is straightforward for any *L*-estimator. Computational results for an option portfolio and a credit-risky bond portfolio demonstrate that the Harrell-Davis estimator, which averages multiple order statistics, is more reliable for risk-management purposes than the sample quantile.

Risk measures are often defined in terms of a quantile of a given distribution. Most popular among these is undoubtedly Value-at-Risk (VaR), which assesses the short-term market risk of a portfolio by identifying the loss that is likely to be exceeded by a specified probability. Thus, the VaR itself is simply a quantile of the portfolio loss distribution. Quantile-based measures are also widely used in portfolio credit-risk applications, where maximum losses and CreditVaR are calculated from distributions of long-term credit losses or counterparty exposures, and in the measurement of earnings-at-risk and operational risk.

While measuring risk accurately remains a critical task for financial institutions, increasing attention is being focused on managing risk, which requires the ability to identify the sources of risk and to assess the impacts of potential trades. Evaluating a quantile's sensitivity to the positions that comprise a portfolio, therefore, is a key requirement for managing not only market risk (VaR), but a variety of risks across the enterprise.

Depending on the application, financial institutions may choose one of several possible methodologies for computing quantiles and their corresponding sensitivities. Techniques for calculating VaR, for example, range from using simple parametric methods, such as the variance-covariance (delta-normal) approach, to simulating the portfolio over a set of historical or Monte Carlo scenarios (see Smithson 2000 for an overview). This paper considers the problem of computing quantile sensitivities in conjunction with simulation approaches, using a general class of non-parametric quantile estimators known as *L*-estimators.

Since the delta-normal method assumes that losses are normally distributed around zero, it yields simple closed-form expressions for VaR and also for its associated risk-management analytics (Litterman 1996, 1997a, 1997b; Garman 1996, 1997). Notably, the trade risk profile (TRP), which relates VaR to the size of the position in a given instrument, is a smooth, convex curve in this case. Thus, it is straightforward to calculate the marginal VaR (i.e., the first-order partial derivative of VaR with respect to position size), the risk contribution and the best hedge position for each instrument in the portfolio.

Recently, Martin et al. (2001) describe a more general parametric approach for calculating VaR sensitivities that is based on the portfolio's cumulant generating function (KGF) and its derivatives at a saddle-point. The method, which also applies to non-normal distributions, is contingent on the ability to calculate the KGF and its first two derivatives, and on the conditional independence of the portfolio holdings.

The scenario-based approach is less reliant on the simplifying assumptions that underlie parametric methods and, thus, it readily accommodates realistic risk-factor distributions and portfolios that contain nonlinear or credit-risky instruments. In this case, it is necessary to estimate VaR from the simulated portfolio losses, which is essentially a statistical quantile-estimation problem. This problem is typically addressed in one of two ways.

So-called semi-parametric methods compute the quantile by fitting a distribution to the sample losses and then using an appropriate inverse transformation (see, for example, Jorion 1996). Extreme value approaches (Embrechts et al. 1998), which focus explicitly on the tail of the loss distribution, also belong in this category. Hughey (1991) evaluates a variety of these methods for estimating extreme quantiles.

In contrast, non-parametric methods make no explicit distributional assumptions —the quantile is derived directly from the data, which are assumed to be an independent, identically distributed sample from an unknown loss distribution. In this case, a point estimate of the quantile is calculated from the order statistics of the sample (the kth order statistic is the kth smallest value in the sample).

Quantile sensitivities are computed in a manner consistent with that of Gourieroux et al. (2000) who provide general forms of the first and second order derivatives of VaR with respect to portfolio allocation. In particular, an asset's marginal VaR equals its expected per-unit loss, conditional on the portfolio loss being equal to the VaR. (This result holds regardless of the method used to calculate VaR, although the actual form of the derivative varies with the chosen estimator.)

One common, non-parametric quantile estimator is the sample quantile, also known as the upper empirical cumulative distribution function value (UECV). For example, given a sample of 100 losses, the UECV estimates VaR at the 95% level as the 96th order statistic (i.e., the fifth-largest loss). Since it relies on only a single order statistic, the sample quantile can exhibit high variability, which reduces its efficiency.

While the accuracy of the UECV-estimated quantile improves with increasing sample size, this is not necessarily true of the estimated quantile sensitivities. In this case, for example, the marginal VaR for an instrument equals its per-unit loss in the particular scenario for which the portfolio loss equals the VaR. As will be shown in this paper, if an instrument's losses do not relate to those of the portfolio in a consistent manner, then the marginal VaR obtained by the UECV will display high variability, regardless of the sample size.

Consistent with the previous observation, Mausser and Rosen (1998a) find that the UECV may produce a jagged TRP, potentially yielding inaccurate estimates of the marginal VaR, risk contribution and best hedge position. To mitigate these effects, the authors fit a smooth curve to the TRP, and then calculate risk analytics based on this approximation.

One drawback of this approach is that it is sensitive to the range of position sizes spanned by the TRP; changing this range may alter the approximating curve and, in turn, the corresponding risk analytics. Furthermore, since it is necessary to construct a portion of the TRP before smoothing, this method may be computationally expensive when one is interested only in the marginal VaR at the current position. This fact is noted by Hallerbach (1999), who proposes instead to calculate the marginal VaR as a weighted average of the instrument returns (or, alternatively, the per-unit losses) over a range of scenarios. The latter approach is, in fact, consistent with using one of a general class of quantile estimators known as *L*-estimators.

L-estimators compute the quantile as a weighted average of multiple order statistics. Note that the UECV is, itself, an *L*-estimator, albeit one that places the entire weight on a single order statistic. Other well-known examples of *L*-estimators include kernel quantile estimators, which effectively smooth the empirical distribution function by spreading each sample observation over an interval with a so-called kernel (Sheather and Marron 1990). The kernel is, typically, a symmetric-probability density function (such as a Gaussian), centred on the observation, and the size of the interval (i.e., the degree of diffusion of the data point) is determined by a bandwidth parameter.

Attractive features of *L*-estimators include their computational simplicity and asymptotic normality (which facilitates the construction of confidence intervals for the quantile). Of particular significance for risk management is the fact that using multiple order statistics improves the robustness of the estimated quantile sensitivities. *L*-estimators are limited by the fact that the quantile estimate cannot lie outside the range of the order statistics, which may be a concern when estimating extreme quantiles based on small samples. Ridder (1998) provides an empirical comparison of several VaR estimation methods, including *L*-estimators and extreme value models.

This paper shows that *L*-estimators are linear in the sizes of the positions that comprise a portfolio, which gives rise to simple representations of the marginal VaR and risk contributions. In this sense, the paper generalizes the methods presented in Mausser and Rosen (1998a) and Hallerbach (1999).

To demonstrate the benefits of utilizing a range of order statistics in estimating risk analytics, we consider the *L*-estimator due to Harrell and Davis (1982), which has been found to perform well empirically (Parrish 1990, Dielman et al. 1994). It is interesting to note that the Harrell-Davis (HD) estimator is asymptotically a Gaussian kernel estimator, although one whose bandwidth is, in fact, suboptimal (Sheather and Marron 1990). While our ensuing discussion is primarily in terms of VaR as estimated from a portfolio loss distribution (in fact, the terms "VaR" and "quantile" are often used interchangeably), it is important to recognize that the concepts extend naturally to any quantile-based measure that derives from an arbitrary distribution.

This paper is organized as follows. First, we describe *L*-estimators in the context of VaR estimation, and contrast the UECV and HD estimators. The next section derives expressions for the marginal VaR and risk contribution for *L*-estimators in general. We then present computational results comparing the performance of the UECV and HD estimators in assessing the market risk of an option portfolio and the credit risk of a bond portfolio. The final section offers some concluding remarks and suggestions for further research.

L-estimators and VaR

For a given time horizon, the $100\alpha\%$ VaR, denoted $VaR(\alpha)$, is the size of loss that will be exceeded with probability $(1 - \alpha)$. Suppose that the loss incurred by a portfolio during the specified period is given by the random variable L, having some (unknown) cumulative distribution function (cdf) F, so that $Prob(L \le y) = F(y)$. The portfolio's $100\alpha\%$ VaR equals the αth population quantile of L, that is, $VaR(\alpha) = F^{-1}(\alpha)$. In the scenario-based approach, $VaR(\alpha)$ (or equivalently $F^{-1}(\alpha)$) is estimated from an independent sample of simulated losses that is drawn from F. We denote this estimated value by Q_α. (It is assumed throughout that $0 < \alpha < 1$.)

The sample of portfolio losses provides an empirical approximation $\hat{F}$ to the true loss distribution F. Consider a set of S scenarios and suppose, for ease of exposition, that the likelihood of each scenario is $1/S$. Let $L_{(k)}$ denote the kth order statistic of the sample, so that $L_{(1)} \le L_{(2)} \le \ldots \le L_{(S)}$. A common way of defining the empirical cdf for the portfolio losses is

$$\hat{F}(y) = \begin{cases} 0 & \text{if } y < L_{(1)} \\ \frac{k}{S} & \text{if } L_{(k)} \le y < L_{(k+1)} \\ 1 & \text{if } y \ge L_{(S)}. \end{cases}$$

One popular estimator of $VaR(\alpha)$ is the inverse of the empirical cdf:

$$\begin{aligned} Q_\alpha &= \hat{F}^{-1}(\alpha) \\ &= L_{(k)} \quad \text{where } \frac{k-1}{S} < \alpha \le \frac{k}{S}. \end{aligned} \tag{1}$$

The UECV estimator is defined similarly, except that in contrast to Equation 1, the second inequality is strict, rather than the first. Thus, the UECV estimate of $VaR(\alpha)$ is

$$Q_\alpha = L_{(k)} \quad \text{where } \frac{k-1}{S} \le \alpha < \frac{k}{S}. \tag{2}$$

Note that the two estimators above yield different results for certain levels of α. For instance, given 100 equally likely scenarios, the inverse empirical cdf (Equation 1) estimates the 95% VaR as $L_{(95)}$ while the UECV (Equation 2) uses $L_{(96)}$.

Both estimators can be viewed more generally as expressing Q_α as a weighted average of the order statistics

$$Q_\alpha = \sum_{k=1}^{S} w_{\alpha,S,k} L_{(k)}. \tag{3}$$

In the case of the inverse empirical cdf, the weight of the kth order statistic in the sample of size S is

$$w_{\alpha,S,k} = \begin{cases} 1 & \text{if } k = \lceil S\alpha + 1 \rceil - 1 \\ 0 & \text{otherwise,} \end{cases}$$

while the UECV uses

$$w_{\alpha,S,k} = \begin{cases} 1 & \text{if } k = \lfloor S\alpha \rfloor + 1 \\ 0 & \text{otherwise.} \end{cases}$$

A quantile estimator having the general form specified in Equation 3, with weights satisfying

$$\sum_{k=1}^{S} w_{\alpha,S,k} = 1$$

is known as an L-estimator. Thus, both the inverse empirical cdf and the UECV are examples of L-estimators that place the entire weight on a single order statistic.

In general, to improve efficiency, L-estimators calculate Q_α as a weighted average of multiple order statistics (see, for example, Dielman et al. 1994). For example, the HD estimator (Harrell and Davis 1982) is based on the fact that, as the sample size increases, the expected value of order statistic $(S + 1)\alpha$ converges to $F^{-1}(\alpha)$ for $0 < \alpha < 1$. Thus, the HD estimator computes Q_α as $E[L_{((S+1)\alpha)}]$, regardless of the integrality of $(S + 1)\alpha$ (as noted by Sheather and Marron (1990), the HD estimator is, in fact, the bootstrap estimator of $E[L_{((S+1)\alpha)}]$). The resulting weights are

$$\begin{aligned} w_{\alpha,S,k} &= \frac{1}{\beta\{(S+1)\alpha,\ (S+1)(1-\alpha)\}} \\ &\quad \times \int_{(k-1)/S}^{k/S} y^{(S+1)\alpha-1}(1-y)^{(S+1)(1-\alpha)-1} dy \\ &= I_{k/S}\{(S+1)\alpha, (S+1)(1-\alpha)\} \\ &\quad - I_{(k-1)/S}\{(S+1)\alpha, (S+1)(1-\alpha)\}, \end{aligned}$$

where $I_x(a,b)$ is the incomplete beta function.

Figure 1 plots the weights for estimating the 0.95 quantile from a sample of 100 observations. Note that unlike the UECV estimator, which places the total weight on the 96th order statistic in this case, the HD estimator distributes the weights among a range of order statistics. As will be demonstrated shortly, this has an implicit smoothing effect that leads to more robust risk analytics, most notably in terms of the marginal VaR and risk contribution, than those provided by the UECV estimator.

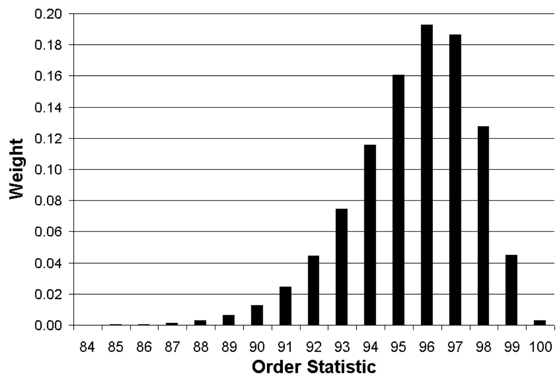

Figure 1: Weights for HD estimator ($S = 100$, $\alpha = 0.95$)

VaR sensitivities with *L*-estimators

Since L-estimators are linear functions of the order statistics, and portfolio losses are linear in the sizes of the positions comprising the portfolio, the resulting VaR estimate is a linear function of the given position sizes. As shown in this section, this fact simplifies the calculation of marginal VaR and risk contributions. Since we now explicitly consider risk to be a function of the sizes of the positions comprising the portfolio, we let $VaR(\alpha; \boldsymbol{x})$ denote the $100\alpha\%$ VaR for the portfolio consisting of positions $\boldsymbol{x}$.

First, it is necessary to express the portfolio loss in terms of the position sizes. Let m_i^o and m_{ij} denote, respectively, the current unit value of instrument i, $i = 1, \ldots, N$, and its Mark-to-Future value in scenario j, $j = 1, \ldots, S$, at the appropriate time horizon. Thus, $\Delta m_{ij} = m_i^o - m_{ij}$ is the unit loss of instrument i in scenario j. If the size of the position in instrument i is x_i, then the loss incurred by an N-instrument portfolio in scenario j is

$$L_j(\boldsymbol{x}) = \sum_{i=1}^{N} \Delta m_{ij} x_i. \tag{4}$$

Let $L_{(k)}(\boldsymbol{x})$ denote the kth order statistic of the set of S sample losses of the portfolio containing the positions $\boldsymbol{x}$. Then, following Equation 3, the $100\alpha\%$ VaR is estimated as

$$VaR(\alpha;\boldsymbol{x}) = \sum_{k=1}^{S} w_{\alpha,S,k} L_{(k)}(\boldsymbol{x}). \tag{5}$$

Note that the weights are independent of the position sizes; they do not need to be re-calculated as the composition of the portfolio changes. Specifically, a set of weights needs only to be calculated once for any given sample size (S) and quantile (α).

To simplify the notation, let $\Delta m_{i(k)}$ denote the per-unit loss of instrument i in the scenario that results in the kth smallest portfolio loss, given the positions $\boldsymbol{x}$ (although we do not explicitly write $\Delta m_{i(k)}$ as a function of $\boldsymbol{x}$, this dependency will be recognized in the subsequent analysis). From Equations 4 and 5, it follows that

$$\begin{aligned} VaR(\alpha;\boldsymbol{x}) &= \sum_{k=1}^{S} w_{\alpha,S,k} \sum_{i=1}^{N} \Delta m_{i(k)} x_i \\ &= \sum_{i=1}^{N} \left\{ \sum_{k=1}^{S} w_{\alpha,S,k} \Delta m_{i(k)} \right\} x_i \\ &= \sum_{i=1}^{N} \omega_{\alpha,S,i}\, x_i, \end{aligned} \tag{6}$$

where

$$\omega_{\alpha,S,i} = \sum_{k=1}^{S} w_{\alpha,S,k} \Delta m_{i(k)}$$

represents the weighted loss per unit of instrument i.

Note that $\omega_{\alpha,S,i}$ is constant as long as the order of the scenarios, ranked by the size of the portfolio loss, does not change. Since this ranking typically changes as the position sizes vary, it follows that $VaR(\alpha; \boldsymbol{x})$ is a piecewise linear function of $\boldsymbol{x}$.

TRP and best hedge position

Recall that the TRP plots VaR against the position size of a particular instrument, assuming that all other positions in the portfolio remain fixed. When VaR is calculated using the delta-normal method, the TRP is a smooth curve with a simple, closed-form representation (see, for example, Mausser and Rosen 1998a). In contrast, the preceding analysis (i.e., Equation 6) shows that under a scenario-based approach, *L*-estimators produce a TRP that is piecewise linear. The smoothness of the resulting TRP depends significantly on the particular choice of *L*-estimator; in particular, an estimator that relies on only a single order statistic, such as UECV, tends to produce a jagged TRP (e.g., Figure 2).

Figure 2: Piecewise linear TRP

The TRP for instrument *i* can be constructed by systematically detecting changes in the order statistics as the position size x_i is varied (such an algorithm, for the UECV estimator, is given in Mausser and Rosen, 1998b). However, this approach, which effectively locates the endpoints of the linear segments comprising the TRP, can be costly from a computational standpoint. Instead, one may choose simply to calculate VaR at a number of different position sizes and then interpolate linearly between them to construct the TRP. Since the intervals can be made arbitrarily small, the TRP can be obtained to any desired level of accuracy.

The best hedge position for instrument *i* is the position size, x_i^*, for which the VaR is as small as possible (i.e., the minimum of the TRP). Given the piecewise linearity of the TRP, identifying the best hedge position requires finding the minimum of all sampled points. As such, obtaining an accurate estimate of the best hedge position is contingent on constructing a TRP that spans an appropriate range of position sizes and on using an interval that is sufficiently small. Since the best hedge position is determined primarily by the overall shape of the TRP, rather than by the slopes of the individual segments, the choice of *L*-estimator tends to have only a minor impact on estimates of x_i^*.

Marginal VaR and risk contribution

From Equation 6 it follows immediately that

$$\frac{\partial\, VaR(\alpha;\boldsymbol{x})}{\partial x_i} = \omega_{\alpha,S,i}, \tag{7}$$

and so $\omega_{\alpha,S,i}$ is the marginal VaR of instrument *i*. The marginal VaR equals the slope of the TRP for instrument *i* at the current position size, and it is, therefore, a piecewise constant function of x_i. Since Equation 7 is not well defined at the endpoints of the segments comprising the TRP, it is necessary to use a modified definition of marginal VaR in this case. At an endpoint x_i', one could, for example, simply compute the marginal VaR at $x_i' \pm \varepsilon$ for suitably small ε, and consider two one-sided derivatives. Note that Equations 6 and 7 imply that the marginal VaR is the expected loss per unit of instrument *i*, given that the portfolio loss equals the VaR, as stated in Gourieroux et al. (2000).

Since VaR is homogeneous of degree one, it satisfies the relation

$$VaR(\alpha;\boldsymbol{x}) = \sum_{i=1}^{N} \frac{\partial\, VaR(\alpha;\boldsymbol{x})}{\partial x_i} x_i$$

and, thus, it admits a marginal decomposition. In general, the marginal contribution of the *i*th position to the portfolio VaR is

$$C(x_i) = \frac{1}{VaR(\alpha;\boldsymbol{x})} \times x_i \frac{\partial\, VaR(\alpha;\boldsymbol{x})}{\partial x_i} \times 100\%$$

which, from Equation 7, can be written as

$$C(x_i) = \frac{\omega_{\alpha,S,i}\, x_i}{VaR(\alpha;\boldsymbol{x})} \times 100\%. \tag{8}$$

Note that the preceding decomposition must be interpreted on a marginal basis; if all positions are scaled by some factor $(1 + \varepsilon)$, where ε is a small constant, then VaR increases by an amount $\varepsilon \times VaR(\alpha; \boldsymbol{x})$ and Equation 8 indicates the relative contribution of the ith instrument to this increase.

Computational results

As noted previously, quantile estimators that rely on a single order statistic (such as UECV) can yield poor risk analytics. This is of particular concern when the performance of a given instrument and that of the overall portfolio are not related in a consistent manner. For example, consider a short straddle position on a stock (i.e., a short call and a short put on the same underlying stock, with identical strike prices and expiration dates), with a strike equal to the current stock price. In this case, large losses can occur when the price of the underlying stock moves either up or down by a significant amount. As a result, two scenarios may yield identical portfolio losses, yet exhibit widely different losses by the individual instruments. For instance, Figure 3 shows that, for a given portfolio loss, the shorted call may incur a small gain if the stock price declines (point A) or a large drop in value (point B) if the stock price increases. Thus, a VaR estimate that is based on only one scenario will result in a marginal VaR that is highly variable.

We now consider two portfolios that exhibit this behaviour, and show how the HD estimator provides more robust risk analytics than the UECV estimator. The first example, discussed previously by Mausser and Rosen (1998a), considers market risk for an option portfolio. The second example, due to Bucay and Rosen (1999), involves analysing credit risk for a portfolio of bonds from emerging markets (see, also, Mausser and Rosen 1999). In the ensuing discussion, it is important to note that the UECV and HD estimators yield similar VaR estimates and best hedge positions; significant differences only become apparent when calculating the associated marginal analytics.

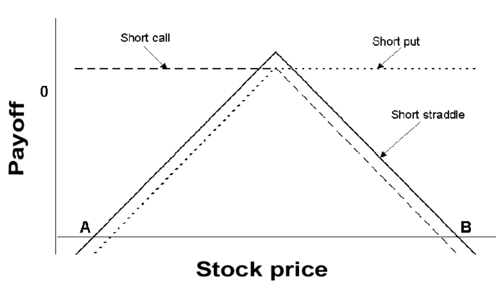

Figure 3: Short straddle pay-off profile

Example: NIKKEI portfolio

Table 1 shows a representative trading desk portfolio, with a current value of 12,493 million JPY, that implements a strangle on two component stocks of the NIKKEI index. In addition to common shares of Komatsu (current price 840,000 JPY) and Mitsubishi (current price 860,000 JPY), the portfolio includes several European call and put options on these equities. Simulating the portfolio over a set of 1,000 Monte Carlo scenarios on the index level (the equity prices are obtained from the *CAPM* model, with both stocks having positive betas) indicates that the loss distribution is far from normal (Figure 4).

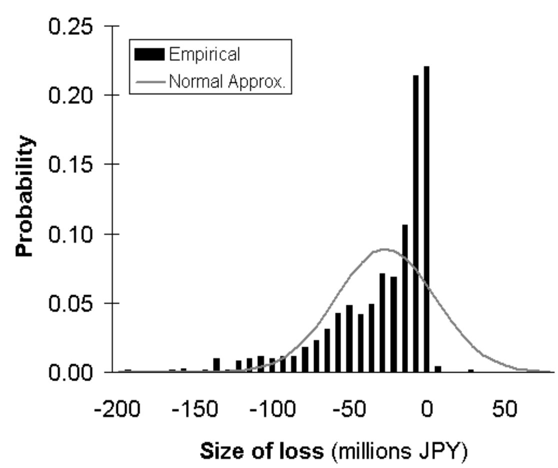

Figure 4: Distribution of losses for the NIKKEI portfolio with best normal approximation (1,000 scenarios)

Instrument	Type	Days to Maturity	Strike Price (10^3 JPY)	Position (x 10^3)	Value (10^3 JPY)
Komatsu	Equity	n/a	n/a	2.5	2,100,000
Mitsubishi	Equity	n/a	n/a	2.0	1,720,000
Komatsu Cjul29 900	Call	7	900	–28.0	–11,593
Mitsubishi Cjul29 800	Call	7	800	–16.0	–967,280
Mitsubishi Csep30 836	Call	70	836	8.0	382,070
Mitsubishi EC 6mo 860	Call	184	860	11.5	563,340
Komatsu Cjun2 760	Call	316	760	7.5	1,020,110
Komatsu Cjun2 670	Call	316	670	22.5	5,150,461
Komatsu Paug31 760	Put	40	760	–10.0	–68,919
Komatsu Paug31 830	Put	40	830	10.0	187,167
Mitsubishi Psep30 800	Put	70	800	40.0	2,418,012

Table 1: NIKKEI portfolio

The VaR(99%) of the portfolio is estimated to be 2.85 million JPY and 2.84 million JPY by the UECV and HD estimators, respectively. In this case, UECV places the entire weight on order statistic 991, while HD distributes the weights among order statistics 969 through 999 (all order statistics outside this range have weights less than 10^{-6}; these weights are set to zero for this analysis).

To appreciate the potential inaccuracies in the risk analytics, it is instructive to first examine a typical TRP. Figure 5 shows the TRP for Mitsubishi EC 6mo 860, one of the call options in the portfolio, constructed using the UECV estimator. The TRP displays an undesirable "jaggedness" in this case, which can potentially result in a marginal VaR (via Equation 7) that is of the wrong sign. Observe, for example, that the TRP slopes upwards, and so the marginal VaR is positive, for positions in the range [5100, 5500] and [6500, 7600], which is inconsistent with the overall shape of the TRP.

To mitigate this problem, Mausser and Rosen (1998a) use a smooth approximation to the TRP and calculate VaR analytics based on this curve. Figure 6 displays one such approximation, obtained by fitting a third-order polynomial to the data in a least squares sense. Note that finding a smooth approximation in this manner requires that a portion of the TRP first be constructed, even if the marginal VaR is desired only for the current position.

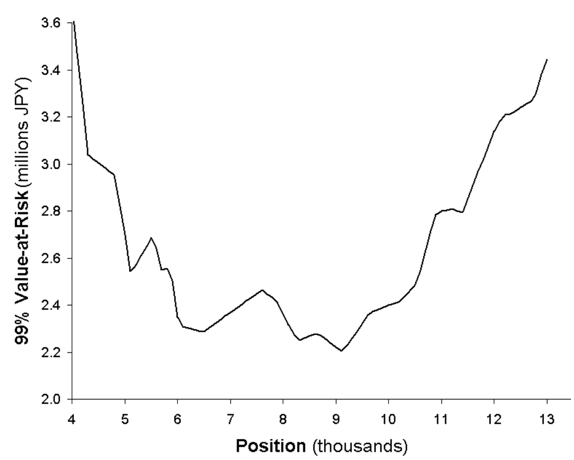

Figure 5: UECV TRP for Mitsubishi EC 6mo 860

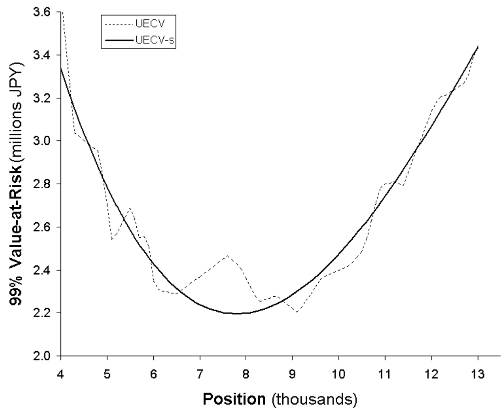

Figure 6: Smoothed UECV TRP for Mitsubishi EC 6mo 860

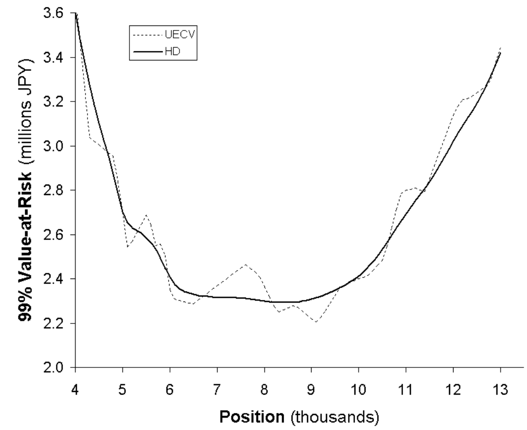

Figure 7: HD and UECV TRP for Mitsubishi EC 6mo 860

	VaR Contribution (%)			Marginal VaR (JPY)		
Instrument	UECV	UECV-s	HD	UECV	UECV-s	HD
Komatsu Cjun2 670	2151	1094	1341	2727	1028	1692
Komatsu Cjun2 760	678	344	424	2576	970	1604
Mitsubishi Csep30 836	477	247	289	1699	653	1026
Mitsubishi EC 6mo 860	232	179	132	575	329	325
Komatsu	202	101	122	2300	857	1389
Mitsubishi	149	75	90	2119	790	1280
Komatsu Paug31 760	53	28	32	–152	–60	–91
Komatsu Cjul29 900	–51	–26	–31	52	20	31
Komatsu Paug31 830	–237	–119	–143	–675	–252	–407
Mitsubishi Cjul29 800	–1166	–591	–706	2078	781	1254
Mitsubishi Psep30 800	–2387	–1232	–1449	–1702	–651	–1029

Table 2: VaR contribution and marginal VaR

The TRP for Mitsubushi EC 6mo 860 obtained using the HD estimator (Figure 7) is less jagged than that of the UECV estimator; the fact that VaR is calculated as a weighted combination of the losses across multiple scenarios has an inherent smoothing effect in this case. Since there is no need to approximate the TRP, it is possible to obtain a robust estimate of the marginal VaR at the current position without having to first construct the TRP.

Table 2 lists the risk contribution and marginal VaR for all positions in the portfolio, as calculated by the UECV and HD estimators and a third-order polynomial approximation to the UECV TRP (UECV-s). As might be expected, the values obtained from the UECV estimator tend to be more extreme than those of the other two approaches, that is, the UECV estimator apparently overestimates the magnitude of the marginal VaR and the VaR contribution.

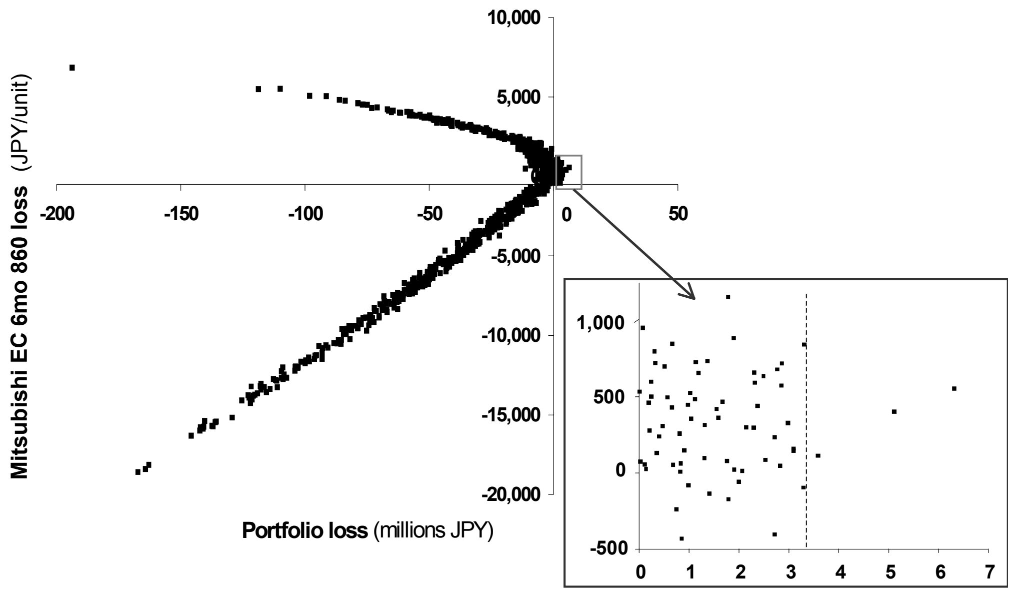

Figure 8: Portfolio losses and Mitsubishi EC 6mo 860 unit losses

However, the signs and relative sizes of the risk analytics are consistent across the three approaches. Note that at the current position sizes, positive risk contributions are due to the long calls, short put and common stocks, all of which appreciate in value when the broader index gains (under the assumed model), while the opposite is true of the short calls and the long puts.

The magnitudes of the VaR contributions are large; since the portfolio is highly leveraged and well hedged, the risks incurred by individual positions tend to offset each other to a large extent. Given the strangle combination, the portfolio loses value for small positive or negative changes in the index. This is evident in Figure 8, which plots the portfolio losses against the per-unit losses for Mitsubishi EC 6mo 860 (the portfolio loses when the option value remains more or less constant). Furthermore, when portfolio losses do occur, they are not highly correlated with those of individual instruments. For example, two scenarios with virtually identical portfolio losses of 3.31 million JPY and 3.29 million JPY correspond to unit losses of 844 JPY and –96 JPY, respectively, for Mitsubishi EC 6mo 860. This suggests potential instability, across samples, in the marginal VaR and risk-contribution estimates that are based on a single order statistic.

Table 3 lists the best hedge positions and the corresponding reductions in VaR that can be achieved. In contrast to the marginal analytics reported in Table 2, the differences between estimators are negligible in this case; the jaggedness of the TRP has minimal impact on the best hedge position, which depends primarily on the overall shape of the curve rather than on its slope. The results suggest that risk can be reduced by taking a greater short position in the index by selling calls and stock or buying puts.

Example: Emerging markets bond portfolio

Portfolio credit risk models account for correlations among the credit transitions of different obligors, thereby recognizing the benefits of diversification as a means of reducing credit risk. Thus, they provide a foundation for applying tools such as marginal VaR (note that the term “maximum loss” is often used in place of “VaR”

Instrument	Position (x 10^3)	Best Hedge Position (x 10^3)			VaR Reduction (%)		
		UECV	UECV-s	HD	UECV	UECV-s	HD
Komatsu Cjun2 670	22.5	20.1	20.0	19.6	41.3	43.0	41.2
Komatsu Cjun2 760	7.5	5.0	4.8	4.4	42.2	44.2	42.2
Mitsubishi Csep30 836	8.0	3.9	4.6	4.2	34.8	35.3	34.2
Mitsubishi EC 6mo 860	11.5	9.1	7.8	8.5	22.9	23.0	19.3
Komatsu	2.5	-0.4	–0.1	–0.3	35.1	35.0	33.7
Mitsubishi	2.0	–1.2	-0.8	–1.0	35.1	35.0	33.7
Komatsu Paug31 760	–10.0	33.2	28.7	33.2	35.2	35.7	34.2
Komatsu Cjul29 900	–28.0	–150.5	–121.0	–145.5	34.7	35.2	33.3
Komatsu Paug31 830	10.0	19.8	19.0	19.8	34.6	35.4	34.0
Mitsubishi Cjul29 800	–16.0	–19.3	–18.8	–19.1	34.8	35.0	33.6
Mitsubishi Psep30 800	40.0	44.2	43.5	43.8	34.7	35.4	34.4

Table 3: Best hedge position and VaR reduction

when dealing with credit risk), risk contributions and best hedges to the problem of managing portfolio credit risk. However, such models may involve simulations that increase the chance of erroneous risk analytics when quantile estimates are based on a single order statistic. Estimators that calculate a weighted average of multiple order statistics, such as the HD estimator, are particularly useful in this case.

In measuring the credit risk of a portfolio of emerging markets bonds, Bucay and Rosen (1999) employ a CreditMetrics methodology (J.P. Morgan 1997) that considers both credit migration and default events. Using implied forward rates, the model first computes exposures, under all possible credit states, for each obligor at a given time horizon (in this case, there are seven credit states plus a default state). A simulation is then performed on the joint credit states of all obligors at the horizon, and a portfolio value is obtained in each scenario by summing the exposures corresponding to the respective credit states. Finally, the portfolio loss is calculated by subtracting the resulting portfolio value from the forward value of the portfolio in the absence of any credit event.

Note that under this model, a given obligor can incur only a small number of possible losses in each scenario (one for each credit state). A quantile estimator that is based on a single order statistic, therefore, will always produce one of these values as the marginal VaR. Since the most likely outcome for each obligor is retaining its existing credit rating, there is a high likelihood of finding a marginal VaR (and also a risk contribution) of zero for a disproportionate number of obligors in this case. Averaging over multiple losses helps to smooth out these sampling effects and, therefore, provides more accurate estimates.

The example portfolio contains 197 long-dated corporate and sovereign bonds, issued by 86 obligors in 29 countries. The mark-to-market value of the portfolio is 8.3 billion USD, and the duration is approximately five years. A simulation of the bond portfolio over 20,000 scenarios and a one-year horizon yields a characteristic credit loss distribution (Figure 9) that is highly skewed and fat tailed.

The VaR(99%) for the bond portfolio is found to be 1.026 billion USD and 1.028 billion USD by the UECV and HD estimators, respectively. In this case, the UECV places the entire weight on order statistic 19,801, while the HD estimator distributes the weights (again, excluding those smaller than 10^{-6}) among order statistics 19,732 to 19,858.

Examining the TRPs for Mexican debt (Figure 10) indicates that the UECV estimator is likely to yield inaccurate risk analytics for this portfolio. In contrast, the HD estimator produces a TRP that better matches the smooth approximation (UECV-s), suggesting that it is more reliable for risk-management purposes.

Table 4 lists the risk contributions and the one-dollar marginal VaR (i.e., the change in VaR that results from investing one additional dollar in an obligor's debt) for the eight largest contributors to the overall credit risk of the portfolio.

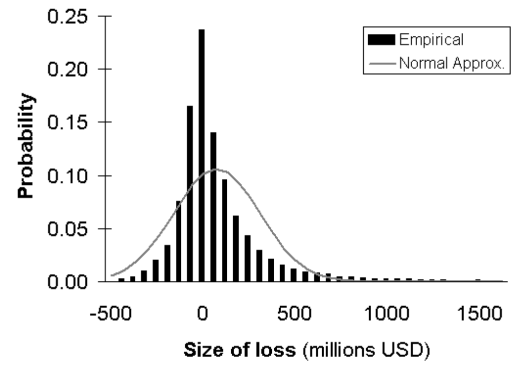

Figure 9: Distribution of losses for the bond portfolio with best normal approximation (20,000 scenarios)

The results of the HD estimator are consistent with those of the smooth approximation (UECV-s). However, the UECV estimator apparently overestimates the marginal VaRs and risk contributions of Peru and Colombia, while underestimating those of Venezuela and Argentina. Furthermore, the UECV estimator produces a marginal VaR and risk contribution of zero for Mexico, exemplifying the problem described previously. The reason for this is readily seen in Figure 11, which shows that the UECV TRP for Mexico is horizontal, not only at the current position of 491 million USD, but also for much of the range between 100 million USD and 800 million USD.

	Current Value	VaR Contribution (%)			$1 Marginal VaR (USD)		
Obligor	(millions USD)	UECV	UECV-s	HD	UECV	UECV-s	HD
Brazil	894	24.69	24.62	25.85	0.28	0.28	0.30
Russia	758	15.12	17.23	19.22	0.20	0.23	0.26
Venezuela	414	7.23	13.96	13.30	0.18	0.35	0.33
Argentina	636	8.83	11.55	12.47	0.14	0.19	0.20
Peru	279	17.54	7.98	7.57	0.65	0.29	0.28
Colombia	608	7.12	2.36	2.35	0.12	0.04	0.04
Mexico	491	0.00	1.70	2.09	0.00	0.04	0.04
Russia CCC	44	2.02	1.70	1.90	0.47	0.39	0.44

Table 4: VaR contribution and marginal VaR

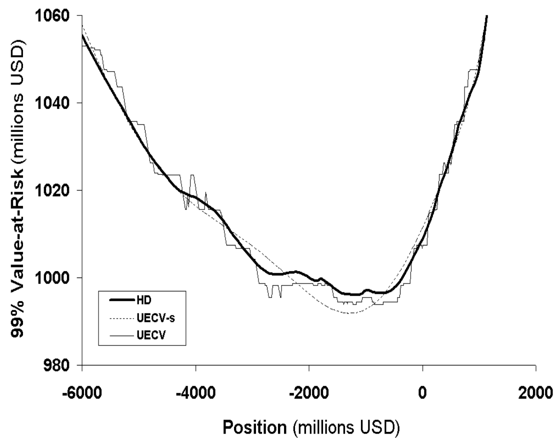

Figure 10: TRPs for Mexico

To further examine the accuracy of the UECV and HD estimators for risk management purposes, we calculate the $1 marginal VaR for Brazil from a sub-sample, selected without replacement, from the full set of 20,000 scenarios. Table 5 reports the mean and standard deviation of the resulting estimates for 1,000 subsamples of various sizes. The HD estimator has the desirable property that, as the sample size increases, the standard deviation of the estimate decreases. In contrast, the variability of the UECV estimates not only exceeds that of the HD estimates, but actually gets larger as the sample size increases in this case.

The reason for this behaviour is readily apparent from the distributions of the Brazilian $1 marginal VaR, as obtained by the two estimators for

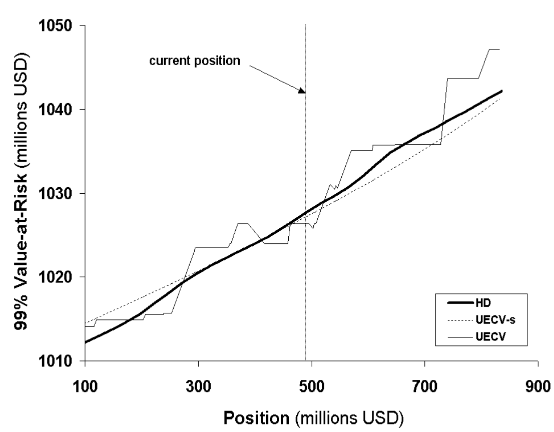

Figure 11: TRPs for Mexico at current position

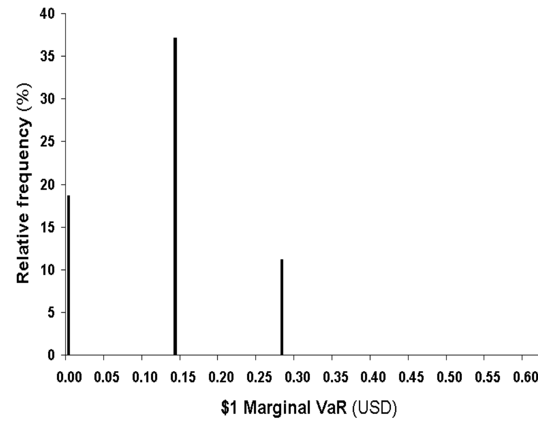

Figure 12: Distribution of $1 marginal VaR for Brazil using UECV estimator (1,000 samples of size 10,000)

samples of size 10,000. The UECV estimator always yields one of four different values (Figure 12), corresponding to four of the eight possible credit states. While the mean value of the estimates is 0.29, their variability is extremely high. In contrast, the HD estimates (Figure 13) are approximately normally distributed around 0.30 (as verified by a Kolmogorov-Smirnov test), and display less variation than the UECV estimates.

The estimators provide similar results in terms of best hedges and the accompanying risk reductions (Table 6). As seen in Figure 10, for example, the overall shapes of the TRPs are, in fact, quite consistent. For this portfolio, minimizing risk requires taking leveraged short positions in the debt of the listed obligors, which might be achieved by entering into a total return swap, for example.

Sample Size	UECV		HD	
1,000	0.28	(0.25)	0.29	(0.08)
5,000	0.31	(0.26)	0.29	(0.04)
10,000	0.29	(0.25)	0.30	(0.03)

Table 5: Mean (standard deviation) of $1 marginal VaR for Brazil based on 1,000 samples

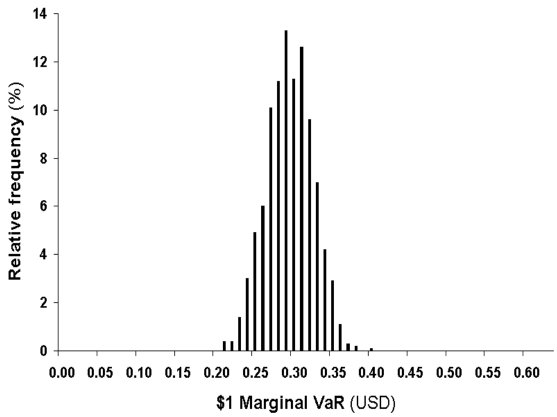

Figure 13: Distribution of $1 marginal VaR for Brazil using HD estimator (1,000 samples of size 10,000)

Conclusions

Scenario-based approaches allow risk to be assessed accurately when parametric methods require unreasonable assumptions, such as for portfolios that contain optionality or that incur credit risk, for example. Given the portfolio losses over a set of scenarios, non-parametric methods for calculating VaR estimate a quantile of the unknown loss distribution directly from the order statistics of the sample losses. *L*-estimators, which compute linear combinations of the order statistics, have desirable computational and statistical properties that make them effective in this regard.

Effective risk management requires not only accurate estimates of the quantile itself, but also of its dependencies on the positions comprising the portfolio. *L*-estimators that are based on only a single order statistic often prove to be inadequate for the latter purpose, even for large samples, since the losses incurred by individual instruments may be poorly correlated with those of the overall portfolio. Thus, estimates of the marginal VaR and risk contribution may exhibit substantial variability in this case.

More robust risk analytics can be obtained from *L*-estimators that compute a weighted average of a range of order statistics since this has an inherent smoothing effect. This paper has shown that the calculation of marginal VaR and risk contribution is straightforward for any *L*-estimator. Computational results for an option portfolio and a credit-risky bond portfolio demonstrate that the Harrell-Davis estimator, which averages multiple order statistics, is more reliable for risk management purposes than the traditional sample quantile (UECV), which considers only one order statistic. The improvement is most evident when calculating marginal VaR and risk contributions, which depend on "local" properties of the TRP (i.e., its slope at a given position size). Differences in estimates of the best hedge

	Current Value	Best Hedge Position (millions USD)			VaR Reduction (%)		
Obligor	(millions USD)	UECV	UECV-s	HD	UECV	UECV-s	HD
Brazil	894	–4,494	–4,077	–4,409	41	41	41
Russia	758	–6,599	–6,687	–5,863	35	34	35
Venezuela	414	–1,375	–1,433	–1,374	34	34	34
Argentina	636	–4,966	–4,843	–4,986	28	28	28
Peru	279	–1,855	–1,741	–1,865	28	28	28
Colombia	608	–27,040	–27,389	–26,960	21	21	21
Mexico	491	–1,392	–1,320	–1,135	3	3	3
Russia CCC	44	–812	–797	–814	27	26	27

Table 6: Best hedge position and VaR reduction

position, which is "global" in the sense that it depends on the overall shape of the TRP, are less significant.

There are several interesting opportunities for extending the results of this paper. Ideally, estimates of the marginal VaR, risk contribution and best hedge position should also include an associated confidence interval. The efficient construction of such confidence intervals would allow risk managers to assess the quality of reported risk analytics better. Furthermore, this paper has considered only the UECV and HD estimators. Additional empirical and theoretical research into the relative merits of various *L*-estimators from a risk-management perspective would allow managers to make more informed choices in this regard.

References

Bucay, N. and D. Rosen, 1999, "Credit risk of an international bond portfolio: A case study," *Algo Research Quarterly,* 2(1): 9–29.

Dielman, T., Lowry, C. and R. Pfaffenberger, 1994, "A comparison of quantile estimators," *Communications in Statistics: Simulation and Computation*, 23(2): 355–371.

Embrechts, P., Resnick, S. and G. Samorodnitsky, 1998, "Living on the edge," *Risk,* 11(1): 96–100.

Garman, M., 1996, "Improving on VaR," *Risk,* 9(5): 61–63.

Garman, M., 1997, "Taking VaR to pieces," *Risk,* 10(10): 70–71.

Gourieroux, C., Laurent, J.P. and O. Scaillet, 2000, "Sensitivity analysis of values at risk," *Journal of Empirical Finance* , 7(3–4): 225–245.

Hallerbach, W.G., 1999, "Decomposing Value-at-Risk: A general analysis," Working Paper, Tinbergen Institute Rotterdam.

Harrell, F.E. and C.E. Davis, 1982, "A new distribution-free quantile estimator," *Biometrika,* 69(3): 635–640.

Hughey, R.L., 1991, "A survey and comparison of methods for estimating extreme right tail-area quantiles," *Communications in Statistics: Theory and Methods,* 20(4): 1463–1496.

Jorion, P., 1996, "Risk^2: Measuring the risk in value at risk," *Financial Analysts Journal,* 52 (November/December): 47–56.

J.P. Morgan, 1997, *CreditMetrics: The Benchmark for Understanding Credit Risk*, Technical Document, 1997, New York, NY: J.P. Morgan Inc.

Litterman, R., 1996, "Hot spots and hedges," *Journal of Portfolio Management,* (Special Issue): 52–75.

Litterman, R., 1997a, "Hot spots and hedges (I)," *Risk,* 10(3): 42–45.

Litterman, R., 1997b, "Hot spots and hedges (II)," *Risk,* 10(5): 38–42.

Martin, R., Thompson, K. and C. Browne, 2001, "VAR: Who contributes and how much?" *Risk,* 14(8): 99–102.

Mausser, H. and D. Rosen, 1998a, "Beyond VaR: From measuring risk to managing risk," *Algo Research Quarterly,* 1(2): 5–20.

Mausser, H. and D. Rosen, 1998b, "Marginal VaR analysis: Parametric and non-parametric methods," Research Paper, Algorithmics Inc.

Mausser, H. and D. Rosen, 1999, "Applying scenario optimization to portfolio credit risk," *Algo Research Quarterly,* 2(2): 19–33.

Parrish, R.S., 1990, "Comparison of quantile estimators in normal sampling," *Biometrics,* 46(1): 247–257.

Ridder, T., 1998, "Basics of statistical VaR-estimation," in *Risk Measurement, Econometrics and Neural Networks: Selected Articles of the 6th Econometric Workshop in Karlsruhe, Germany,* (G. Bol, G. Nakhaiezadeh and K.-H. Vollmer, Eds.), Physica-Verlag, Heidelberg, 161–187.

Sheather, S.J. and J.S. Marron, 1990, "Kernel quantile estimators," *Journal of the American Statistical Association,* 85(410): 410–416.

Smithson, C., 2000, "VAR—a work in progress," *Risk*, 13(10): 49–51.