Beyond VaR: From Measuring Risk to Managing Risk

Helmut Mausser and Dan Rosen

This paper examines tools for managing, as opposed to simply monitoring, a portfolio's Value-at-Risk (VaR). These tools include the calculation of VaR contribution, marginal VaR and trade risk profiles. We first review the parametric, or delta-normal, versions of these tools and then extend them to the simulation-based, or non-parametric, case. We analyze two sample portfolios: one, consisting of foreign exchange contracts, is well-suited for parametric analysis while the other, which contains European options, is best addressed with simulation-based methods. The limitations of the simulation-based approach, due to the potential effects of sampling error, are also discussed.

Financial institutions worldwide have devoted much effort to developing enterprise-wide systems that integrate financial information across their organizations to *measure* their institution's risk. Probabilistic measures, such as Value-at-Risk (VaR), are now widely accepted by both financial institutions and regulators for assigning risk capital and monitoring risk. Since development efforts have been driven largely by regulatory and internal requirements to report risk numbers, tools needed to *understand* and *manage* risk across the enterprise have generally lagged behind those designed to measure it.

Measuring risk is a *passive* activity; simply knowing one's VaR does not provide much guidance for managing risk. In contrast, risk management is a *dynamic* endeavor and it requires tools that help identify and reduce the sources of risk. These tools should lead to an effective utilization of the wealth of financial products available in the markets to obtain the desired risk profiles.

To achieve this, a comprehensive risk manager's toolkit must provide the ability to

- represent complex portfolios simply
- decompose risk by asset and/or risk factor
- understand how new trades affect the portfolio risk
- understand the impact of instruments' non-linearities and of non-normal risk factor distributions on portfolio risks
- understand complex, non-intuitive, market views implicit in the portfolio as well as in the investment policy or market liquidity
- generate potential hedges and optimize portfolios.

Robert Litterman (1996a, 1996b, 1997a, 1997b) recently described a comprehensive set of analytical risk management tools extending some of the insights originally developed by Markowitz (1952) and Sharpe (1964). Developed in close collaboration with the late Fisher Black and his colleagues at Goldman Sachs, these tools are based on a linear approximation of the portfolio to measure its risk and assume a joint (log)normal distribution of the underlying market risk factors, similar to the RiskMetrics VaR methodology (J.P. Morgan 1996). Litterman further emphasized the dangers of managing risk using only such linear approximations. However, in spite of their onerous assumptions, the insights provided by these tools are very powerful and hence constitute a solid basis for a risk management toolkit. (The reader is also referred to the related papers by Mark Garman (1996, 1997) on marginal VaR and risk decomposition.)

This is the first of a series of papers in which we present an *extended simulation-based risk management toolkit* developed on top of the analytical tools presented by Litterman. Simulation-based tools provide additional insights when the portfolio contains non-linearities, when the market distributions are not normal or when there are multiple horizons. In particular, these tools should prove very useful not only for market risk, but also for credit risk, where the exposure and loss distributions are generally skewed and far from normal. We further demonstrate that simulation-based tools can be used, sometimes even more efficiently, with other risk measures in addition to VaR. Indeed, they also uncover limitations of VaR as a coherent risk measure, as has been demonstrated by Artzner et al. (1998).

Simulation-based methods to measure VaR (historical or Monte Carlo) are generally much more computationally intensive than parametric methods (such as the delta-normal method popularized by RiskMetrics). Advances in computational simulation methods and hardware have rendered these methods practical for enterprise-wide risk measurement. However, it is widely believed that risk management tools based on simulation are impractical since they require substantial additional computational work (Dowd 1998). We demonstrate that efficient computational methods are available which generally require little or no additional simulation to obtain risk management analytics.

In this paper, we extend *marginal VaR analysis* to a simulation-based environment, compare the method with the parametric approach and apply it to two practical examples. We demonstrate how one can efficiently obtain the changes in the portfolio VaR, as measured by a simulation, that result from adding a small amount (or percentage) of an asset to the portfolio. We show that, since VaR is a homogeneous function of the positions, we can obtain an additive portfolio decomposition based on marginal VaR, as in the parametric case. We also investigate the trade risk profiles of a single asset or a class of assets. We discuss the properties of these tools, the errors that arise due to sampling, their limitations and possible extensions.

This paper is organized as follows. We first review parametric VaR and use its associated risk management tools to analyze a portfolio of foreign exchange forwards. We then derive the simulation-based tools and discuss the potential effects of sampling error. To demonstrate the methodology, we re-examine the foreign exchange portfolio (obtaining results consistent with the parametric version) and also consider a portfolio of stock options for which the parametric approach is inappropriate. We conclude by suggesting directions for further study.

Parametric VaR

The parametric, or delta-normal, method for calculating VaR assumes the existence of a set of market risk factors whose log price changes are joint normally distributed with zero mean; that is, if r_k is the log return on risk factor k, then $\boldsymbol{r} \sim N(0, \mathbf{Q}^*)$, where $\mathbf{Q}^*$ is the covariance matrix of risk factor returns. Consider a portfolio composed of positions $\boldsymbol{x}$, where x_i is the size of the holding in instrument i, for $i = 1, 2, ..., N$. As shown in the Appendix, the portfolio's $100(1-\alpha)\%$ VaR (which we denote $VaR(\boldsymbol{x})$, implicitly recognizing its dependence on α) is

$$VaR(\boldsymbol{x}) = \sqrt{\boldsymbol{m}(\boldsymbol{x})^T \mathbf{Q} \boldsymbol{m}(\boldsymbol{x})} \qquad (1)$$

where $\mathbf{Q} = z_\alpha^2 \mathbf{Q}^*$ is a scaled covariance matrix (z_α denotes the standard normal z-value that delimits a probability of α in the right tail) and $\boldsymbol{m}(\boldsymbol{x})$, known as the **VaR map** of the portfolio, is a vector of the portfolio's exposure to the risk factors. The elements of $\boldsymbol{m}(\boldsymbol{x})$ equal the monetary value of the portfolio's position in each risk factor. Thus, the VaR map provides a reduced, or simplified, view of the portfolio from a risk management perspective. Note that

$$\boldsymbol{m}(\boldsymbol{x}) = \sum_{i=1}^{N} \boldsymbol{m}^i x_i \qquad (2)$$

where vector $\boldsymbol{m}^i$ is the VaR map of one unit of the i-th instrument (i.e., m_k^i is the exposure to risk factor k that results from holding a single unit of instrument i). Equation 2 shows that the portfolio VaR map is the sum of the instruments' VaR maps, weighted by position.

Trade risk profile and best hedge position

Knowledge of how VaR changes with position size is critical for effective risk management. If we plot the portfolio VaR against the size of the position in a given instrument i (all other positions being fixed), we obtain the **trade risk profile** (TRP). As shown in Figure 1, the TRP has a unique minimum, which occurs at the best hedge position, x_i^{bh}. The best hedge position can be found analytically as described in the Appendix.

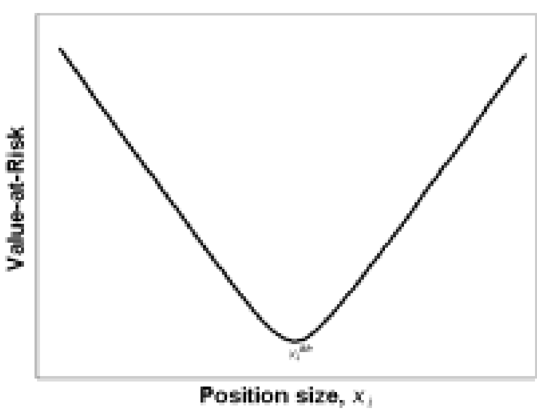

Figure 1: Trade risk profile

Marginal VaR

Managing risk requires an understanding of how new trades affect the portfolio risk. Thus, let us now consider the calculation of the marginal VaR, which measures the impacts of small changes in risk factor exposures or instrument positions on the portfolio VaR. From Equation 1, we find that the VaR gradient with respect to the risk factor exposures is

$$\nabla_{\boldsymbol{m}} VaR(\boldsymbol{x}) = \frac{Q\boldsymbol{m}(\boldsymbol{x})}{VaR(\boldsymbol{x})} \qquad (3)$$

The k-th element of $\nabla_{\boldsymbol{m}} VaR(\boldsymbol{x})$ is the change in VaR that results from increasing the portfolio's exposure to the k-th risk factor (i.e., $m_k(\boldsymbol{x})$) by a single monetary unit.

Since the VaR map of the portfolio is the sum of the VaR maps for the positions (Equation 2), it follows that the derivative of VaR with respect to the i-th position is

$$\frac{\partial VaR(\boldsymbol{x})}{\partial x_i} = (\boldsymbol{m}^i)^T (\nabla_{\boldsymbol{m}} VaR(\boldsymbol{x})) \qquad (4)$$

Equation 4 indicates the change in VaR due to adding one unit of instrument i to the portfolio (if $x_i < 0$, then this corresponds to reducing the short position). Note that this is simply the derivative of the trade risk profile for instrument i at the current position x_i.

VaR contribution

By decomposing VaR, a risk manager is able to target the most significant sources of risk, or the portfolio's so-called "Hot Spots". This task is complicated by the fact that VaR is a sub-additive measure: the portfolio VaR is typically less than the sum of the individual position VaRs. However, since VaR is a homogeneous function (i.e., $VaR(a\boldsymbol{x}) = a \cdot VaR(\boldsymbol{x})$), it admits a marginal decomposition. Note that if we multiply Equation 4 by the position and sum over all holdings in the portfolio, we obtain

$$\sum_{i=1}^{N} x_i \frac{\partial VaR(\boldsymbol{x})}{\partial x_i} = VaR(\boldsymbol{x}) \qquad (5)$$

In Equation 5, each term in the summation is the product of position size and the rate of change of VaR with respect to that position. This essentially represents the rate of change of VaR with respect to a small *percentage* change in the size of the position. Let us define

Instrument	Currency	Days to Maturity	Strike Price (USD)	Position (x 10^6)	Value (x 10^3 USD)
CAD/USD .73 100d	CAD	100	0.73	0.5	2.5
CAD/USD .74 30d	CAD	30	0.74	1.0	–8.3
DEM/USD .57 60d	DEM	60	0.57	6.0	73.2
DEM/USD .59 120d	DEM	120	0.59	5.0	–28.2
FRF/USD .16 40d	FRF	40	0.16	8.0	83.3
JPY/USD .0091 11d	JPY	11	0.0091	10.0	–0.9

Table 1: FX portfolio

$$C(x_i) = \frac{1}{VaR(\boldsymbol{x})} \times x_i \frac{\partial VaR(\boldsymbol{x})}{\partial x_i} \times 100\% \qquad (6)$$

to be the percentage contribution to VaR of the i-th position. Equation 6 must be interpreted on a marginal basis; it indicates the relative contributions to the change in VaR that results if all positions are scaled by the same amount. Note that at the best hedge position, a position's marginal VaR, and therefore also its VaR contribution, is zero.

Similarly, multiplying both sides of Equation 3 by $\boldsymbol{m}(\boldsymbol{x})^{\mathrm{T}}$ shows that VaR is equal to the inner product of the VaR map and the VaR gradient with respect to the risk factor exposures. We can therefore define

$$C(m_k(\boldsymbol{x})) = \frac{1}{VaR(\boldsymbol{x})} \times m_k(\boldsymbol{x}) \frac{\partial VaR(\boldsymbol{x})}{\partial m_k(\boldsymbol{x})} \times 100\% \qquad (7)$$

to be the percentage contribution to VaR of the k-th risk factor. Again, Equation 7 must be interpreted on a marginal basis.

An example FX portfolio

Table 1 shows a portfolio of foreign exchange (FX) forward contracts as of July 1, 1997. Suppose that the exchange rates, in USD, are 0.73 (CAD), 0.58 (DEM), 0.17 (FRF) and 0.0090 (JPY). The total value of the portfolio is 122,000 USD and its one-day 99% VaR is 78,000 USD.

For this example, we elect to use the RiskMetrics risk factor data set for computing the parametric VaR. The portfolio's VaR map, along with the marginal VaR and VaR contribution of each risk factor, are shown in Table 2. The risk factors in Table 2 are sorted in order of decreasing VaR contribution. Since the forwards are contracts to purchase foreign currency with USD, the VaR map consists of long positions in FX spots and foreign interest rates, and short positions in US interest rates. The magnitudes of these positions indicate that the portfolio has significant exposure to the DEM/USD exchange rate and 30- and 90-day interest rates in the US and Germany. Note that exposure to any risk factor can be eliminated by undoing the corresponding VaR map position (e.g., Canadian currency risk can be removed by selling 1.09 million USD worth of Canadian dollars).

The primary source of portfolio risk is currently the DEM exposure, as indicated by the fact that it contributes 83% of the VaR. Conversely, the CAD exposure is actually acting as a hedge for the portfolio, as evidenced by its negative VaR contribution (–0.75%). Furthermore, the VaR gradient indicates that, at the margin, the VaR can be reduced by 5.41 USD for every additional USD of exposure to the Canadian dollar. This may be somewhat surprising given that the portfolio is currently long CAD; it is a useful illustration of the fact that portfolio risk depends not only on the individual risk factors themselves, but also on their correlation (in this case, the Canadian dollar is negatively correlated with the other currencies).

The previous analysis considers the effects of general market factors on the VaR. The analysis at the position level (Table 3) yields consistent

Risk Factor	VaR Map $m_k(\mathbf{x})$ (x 10^3 USD)	Marginal VaR $\frac{\partial VaR(\mathbf{x})}{\partial m_k(\mathbf{x})}$ (x 10^{-4} USD)	VaR Contribution $C(m_k(\mathbf{x}))$ (%)
DEM/USD exchange	6,327	102.67	82.99
FRF/USD exchange	1,355	95.62	16.55
JPY/USD exchange	90	97.51	1.12
US 90-day rate	–4,146	–0.10	0.06
Germany 90-day rate	3,698	0.04	0.02
Germany 180-day rate	602	0.21	0.02
Canada 30-day rate	728	0.13	0.01
US 180-day rate	–628	–0.08	0.01
US 30-day rate	–5,631	0.01	0.00
Germany 30-day rate	3,517	–0.01	0.00
Canada 90-day rate	358	0.09	0.00
France 90-day rate	98	–0.21	0.00
Japan 30-day rate	33	0.01	0.00
Canada 180-day rate	20	–0.83	0.00
France 30-day rate	1,511	–0.03	–0.01
CAD/USD exchange	1,090	–5.41	–0.75

Table 2: Risk factor data for the FX portfolio (ranked by VaR contribution)

results. Together, the two DEM contracts contribute approximately 83% of the current portfolio risk while the CAD contracts, with a total contribution of –0.70%, act as a hedge. Note that these values agree well with the VaR contributions of the DEM (83%) and CAD (–0.75%) risk factor exposures. The marginal VaRs indicate that increasing the positions in the DEM, FRF and JPY contracts results in greater portfolio risk while a similar increase in the CAD contracts reduces risk. This is also reflected by the best hedge positions; for the DEM, FRF and JPY contracts, the best hedges are smaller than the current positions (and in fact suggest shorting contracts), but they are larger than the current positions in the case of the CAD contracts.

The impact of holding the best hedge position in a given instrument can be measured in terms of the percentage reduction in VaR that can be achieved (i.e., the resulting decrease in VaR expressed as a percentage of the current VaR). At their best hedge positions, the DEM contracts each reduce the VaR by almost 88% while each CAD contract offers a much smaller reduction of only 0.2%. In many cases, however, it may simply not be feasible to hold an instrument at its best hedge position. For example, being short seven million units of DEM/USD .57 60d contracts may well run counter to the underlying objectives of the portfolio. Thus, it is often useful to consult the trade risk profile (e.g., Figure 2) to determine the VaR reduction that can be achieved within practical limitations.

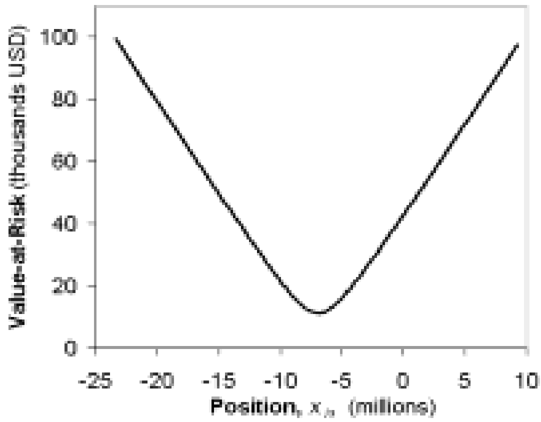

Figure 2: Trade risk profile for DEM/USD .57 60d

Simulation-based VaR

The simulation-based approach to VaR calculation relies on a complete valuation of the portfolio under a set of scenarios, which may derive from historical data or a Monte Carlo simulation. Given a particular "base case" scenario (e.g., representative of current market conditions), it is straightforward to calculate the gain or loss in portfolio value in each scenario. Let v_i^o denote the unit value of instrument i in the base scenario and v_{ij}^t denote its unit value in scenario j at some future time t. We refer to v_{ij}^t as a **mark-to-future** value for instrument i. Since we assume exclusively a one-day time horizon for calculating VaR, we will hereafter dispense with the t superscript to improve readability. Let us define $\Delta v_{ij} = v_i^o - v_{ij}$ to be the unit loss of instrument i in scenario j. If the current position in instrument i is x_i, then the loss (note that a gain is a negative loss) incurred by the portfolio in scenario j is

$$L_j(\mathbf{x}) = \sum_{i=1}^{N} x_i \, \Delta v_{ij} \tag{8}$$

Suppose that the likelihood, or weight, of scenario j is p_j. If we order the losses from largest to smallest (since losses can be negative when the portfolio gains in value, "smallest" is taken here to mean "most negative") and calculate the cumulative scenario probability, then the non-parametric $100(1-\alpha)\%$ VaR, or nVaR, equals the loss in that scenario for which the cumulative probability first meets or exceeds α. We refer to this scenario as the **threshold scenario**. To simplify the notation, we denote the threshold scenario simply as s^o, implicitly recognizing its dependence on $\mathbf{x}$ and α.

For example, consider a portfolio that is evaluated over a set of 100 scenarios. Table 4 shows the five largest losses, in decreasing order of magnitude, along with their respective

Instrument	Marginal VaR $\frac{\partial VaR(x)}{\partial x_i}$ (x 10^{-4} USD)	VaR Contribution $C(x_i)$ (%)	Current Position x_i (x 10^6)	Best Hedge Position x_i^{bh} (x 10^6)	VaR Reduction (%)
DEM/USD .57 60d	59.24	45.4	6.0	–7.0	87.7
DEM/USD .59 120d	58.97	37.7	5.0	–8.0	87.6
FRF/USD .16 40d	16.19	16.5	8.0	–38.3	79.2
JPY/USD .0091 11d	0.88	1.1	10.0	–209.2	13.2
CAD/USD .73 100d	–3.80	–0.2	0.5	1.4	0.2
CAD/USD .74 30d	–3.85	–0.5	1.0	1.9	0.2

Table 3: Instrument data for the FX portfolio (ranked by VaR contribution)

Scenario Number	Loss	Probability	Cumulative Probability
27	10,000	0.010	0.010
82	9,500	0.030	0.040
50	8,800	0.010	0.050
11	8,600	0.020	0.070
63	8,100	0.005	0.075

Table 4: Simulation-based VaR example

scenarios, probabilities and cumulative probabilities. For this particular portfolio and scenario set, the 95% and 98% nVaRs are 8,800 and 9,500, respectively. Note that in the latter case, $p_{27} + p_{82} = 0.04 > 0.02$ and so one might argue that some value between 10,000 and 9,500 provides a better estimate for the 98% nVaR. A discussion of the merits of such interpolation schemes is beyond the scope of this paper; we simply note that the above approach may yield a smaller nVaR, relative to interpolated values, in some cases.

While valuing the portfolio under large numbers of scenarios can be a computationally intensive task, nVaR analysis has the desirable property of requiring only a single simulation. Once the instruments' mark-to-future values have been obtained, Equation 8 can be used to calculate losses for individual holdings (and hence the portfolio) under subsequent changes in the positions.

Trade risk profile and best hedge position

Recall that the trade risk profile plots the level of risk (VaR or nVaR) against the position taken in a particular instrument. In the parametric case, the resulting curve is smooth and has a unique minimum at the best hedge position. The nVaR's dependency on a finite number of scenarios implies that the non-parametric trade risk profile (nTRP) is piecewise linear (Figure 3). As shown in the Appendix, the nTRP consists of multiple segments, each corresponding to a threshold scenario that is in effect for a given range of positions. Unlike the parametric case, the nTRP may have multiple local minima, and therefore, finding the best hedge position (i.e., the global minimum), x_i^{nbh}, may require considerable computational effort. However, this effort lies only in tracing out the nTRP, rather than re-pricing the instruments, and so computational time is typically far less than that required for a full simulation.

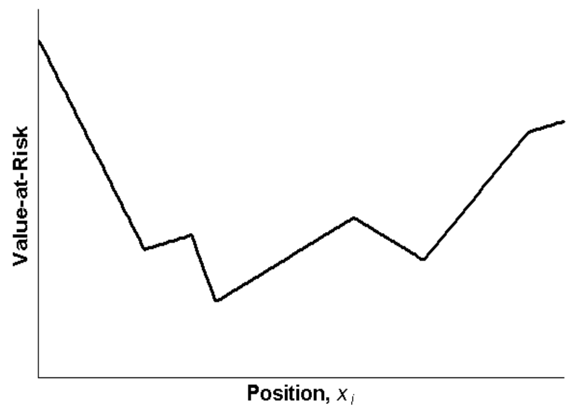

Figure 3: Simulation-based trade risk profile

Marginal nVaR

One might anticipate that calculating the marginal nVaR involves making a small positional change, re-simulating the portfolio and recalculating nVaR. Fortunately, this is not required. From the definition of nVaR as the loss in the threshold scenario

$$nVaR(\boldsymbol{x}) = \sum_{i=1}^{N} x_i \, \Delta v_{is^o} \tag{9}$$

it follows that nVaR is linear in $\boldsymbol{x}$. Let us assume for the moment that the threshold scenario remains unchanged for small variations in the positions. In this case, the derivative of nVaR with respect to the i-th position is

$$\frac{\partial\, nVaR(\boldsymbol{x})}{\partial x_i} = \Delta v_{is^o} \tag{10}$$

Thus, the i-th component of the nVaR gradient is simply the difference between the instrument's values in the base and threshold scenarios.

Let us now examine the marginal nVaR in light of the piecewise linearity of the nTRP. In doing so, we will make reference to Figure 4, which illustrates two adjacent segments of a nTRP for some instrument i. In this case, positions (x_i) of 100, 200 and 300 result in portfolio nVaR values of 25,000, 10,000 and 15,000, respectively. The slope of the first segment is –150 while that of the second segment is 50. Recall that positions in all instruments other than i are held fixed.

Since the nTRP is piecewise linear, all positions in instrument i that lie on the same segment have an identical gradient, whose i-th component is simply the slope of that segment; that is, for all $100 < x_i < 200$, increasing (decreasing) the position in instrument i by a sufficiently small amount δ decreases (increases) the portfolio nVaR by 150δ. In this case, a "sufficiently small" change in position can be calculated precisely as a decrease of up to $x_i - 100$ or an increase of up to $200 - x_i$.

Now consider a position of 200 in instrument i, which corresponds to a change in the threshold scenario. As is evident in Figure 4, the gradient is not well-defined at this point; its i-th component changes abruptly from –150 to 50. To deal with this lack of continuity in the gradient at such points, it is necessary to consider two one-sided sub-gradients. It should be apparent, however, that knowledge of the slopes and endpoints of the segments comprising the nTRP allows marginal nVaR information, and the range for which it is valid, to be reported for any position.

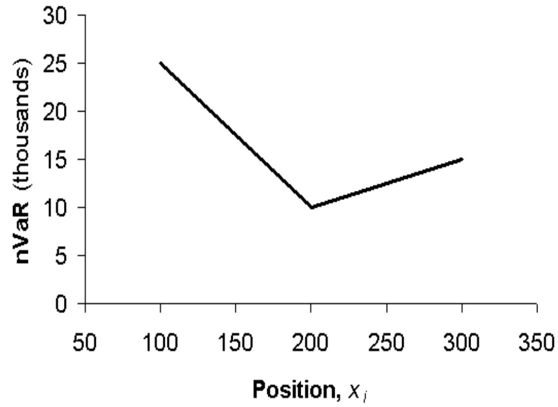

Figure 4: Two segments of a nTRP

nVaR contribution

Equation 9 immediately provides us with a risk decomposition – nVaR is the sum of the position losses in the threshold scenario. Thus, the percentage contribution to nVaR of the i-th position is simply

$$nC(x_i) = \frac{1}{nVaR(\boldsymbol{x})} \times x_i \Delta v_{is^o} \times 100\% \tag{11}$$

Note that Equation 11 is identical to Equation 6 in that the risk contribution is based on the product of the position and the marginal nVaR. Thus, as in the parametric case, the above decomposition must be interpreted on a marginal basis. If we scale all positions by some factor $(1 + \varepsilon)$, where ε is a small constant, then nVaR increases by an amount $\varepsilon \times nVaR(\boldsymbol{x})$ and Equation 11 indicates the relative contribution of the i-th instrument to this increase.

Implementation considerations

The simulation-based approach to VaR, as described in this paper, depends entirely on the scenarios used in the simulation; changing the scenarios is likely to yield different values for nVaR as well as for the related marginal risk measures. Thus, it is important to recognize the possible effects of sampling error on the reported values. In particular, while increasing the number of scenarios generally improves the reliability of nVaR as an estimate of the true Value-at-Risk, the marginal nVaR (and similarly, the nVaR contribution) may still exhibit considerable

Instrument	Position	Scenario #1 mark-to-future values	Scenario #2 mark-to-future values
1	100	10	5
2	50	10	20
Portfolio		1,500	1,500

Table 5: Sample two-instrument portfolio

variability as more scenarios are sampled. To improve the accuracy of these values, we propose using a smooth approximation to the nTRP.

The problem: sensitivity to the threshold scenario

Recall from Equation 10 that the marginal nVaR is determined exclusively by an instrument's values in the base and threshold scenarios. Obtaining consistent marginal nVaR estimates, then, requires that scenarios resulting in similar losses also have similar mark-to-future values for the instruments. However, this may not hold in practice, as illustrated by the following example. Consider a portfolio consisting of only two positions that is simulated over two scenarios (Table 5). The portfolio has an identical value (and therefore an identical loss) in both scenarios, yet the instruments' mark-to-future values (and the marginal nVaRs) are quite different. Therefore, the marginal nVaR is extremely sensitive to the threshold scenario.

Since an instrument's marginal nVaR equals the slope of the nTRP for that instrument at the current position, it follows that adjacent segments of the nTRP can have markedly different slopes (recall from Figure 4 that adjacent segments of the nTRP meet at points where the respective threshold scenarios incur the same loss). Increasing the number of scenarios tends to shorten the average length of segments comprising the nTRP. However, adjacent segments do not necessarily become better "aligned" in the sense of having similar slopes. Thus, the marginal nVaR may not exhibit the convergence that might be expected as more scenarios are sampled.

The solution: smoothing the nTRP

Fitting a smooth curve to the nTRP, and then taking the derivative of this curve, tends to provide a more robust estimate of the true marginal VaR. Essentially, this approach removes the "noise" that is present in the nTRP.

One might consider using splines or fitting a polynomial function, $P_i(x_i)$, to the nTRP of instrument i (specifically, to the endpoints of the segments) in the least squares sense. Clearly, the degree of $P_i(x_i)$ should be chosen so that the curve provides a reasonable approximation to the nTRP, without over-fitting the points. A visual comparison of the two curves will generally establish the suitability of $P_i(x_i)$ for the range of positions being considered. From this approximation, one can then obtain the following estimates:

- marginal nVaR at position x_i

$$\frac{\partial\, nVaR(\boldsymbol{x})}{\partial x_i} \approx P_i'(x_i)$$

- best hedge position

$$x_i^{nbh} \approx \left\{ x_i^* \middle|\ P_i(x_i^*) \le P_i(x_i) \textit{ for all } x_i \textit{ in the range of interest} \right\}$$

- nVaR contribution at position $\boldsymbol{x}$

$$nC(x_i) \approx \frac{x_i\, P_i'(x_i)}{\sum_{l=1}^{N} x_l\, P_l'(x_l)} \times 100\% \qquad (12)$$

The FX portfolio revisited

To compare the parametric and simulation-based VaR analyses, the FX portfolio was simulated over a set of 1,000 Monte Carlo scenarios. The one-day 99% nVaR is 77,000 USD, which differs from the parametric value by less than 2%. The resulting loss histogram is approximated well by a normal distribution (Figure 5) and the nVaR analysis (Table 6) is consistent with its parametric counterpart in Table 3. The final column of Table 6 gives the range for which the marginal nVaR remains unchanged (i.e., the "length" of the current segment of the nTRP for each instrument). Note the close agreement between the parametric and simulation-based trade risk profiles for DEM/USD .57 60d (Figure 6).

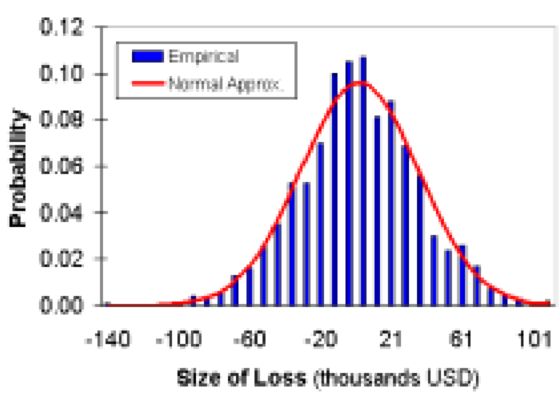

Figure 5: Distribution of losses for the FX portfolio with best normal approximation (1,000 scenarios)

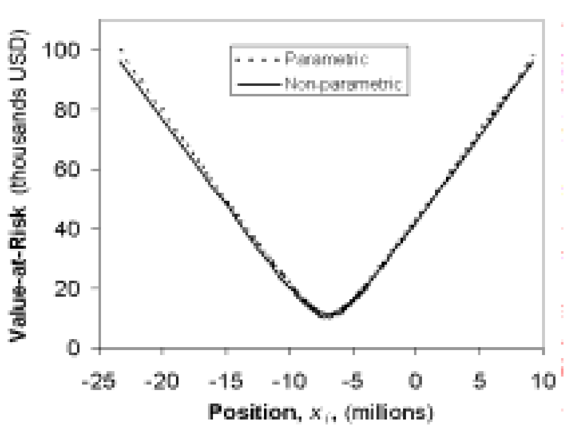

Figure 6: TRP and nTRP for DEM/USD .57 60d

Example - the NIKKEI portfolio

Table 7 shows a portfolio that implements a butterfly spread on the NIKKEI index, as of July 1, 1997. 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. The total value of the portfolio is 12,493 million JPY and its parametric one-day 99% VaR is 115 million JPY.

This portfolio, which may be representative of the positions held by a trading desk, makes extensive use of options to achieve the desired payoff profile. A histogram showing the distribution of losses over a set of 1,000 Monte Carlo scenarios (Figure 7) indicates that the normal distribution fits the data poorly, and that the parametric VaR is likely to over-estimate the true Value-at-Risk. Indeed, simulating the portfolio over these 1,000 scenarios results in a one-day 99% nVaR of 2.9 million JPY, reflecting the fact that the portfolio is well-hedged. Because the parametric VaR measures the risk poorly in this case, we perform only a simulation-based analysis.

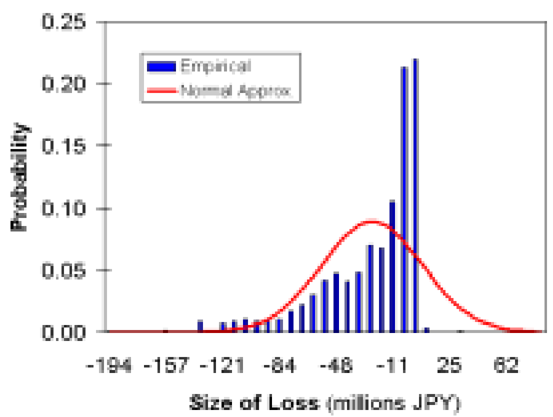

Figure 7: Distribution of losses for the NIKKEI portfolio with best normal approximation (1,000 scenarios)

We compute the contributions, marginal nVaRs and best hedge positions using a pure non-parametric approach (i.e., corresponding to the piecewise-linear nTRP), as well as a third-order

Instrument	Marginal nVaR $\frac{\partial\, nVaR(x)}{\partial x_i}$ (x 10^{-4} USD)	nVaR Contribution $nC(x_i)$ (%)	Current Position x_i (x 10^6)	Best Hedge Position x_i^{nbh} (x 10^6)	nVaR Reduction (%)	Valid Range for Marginal nVaR (x 10^6)
DEM/USD .57 60d	59.73	46.6	6.0	–7.4	87.0	[2.7, 9.2]
DEM/USD .59 120d	59.62	38.7	5.0	–8.2	86.9	[1.7, 8.2]
FRF/USD .16 40d	15.03	15.6	8.0	–38.0	78.6	[3.4, 20.8]
JPY/USD .0091 11d	1.01	1.3	10.0	–258.6	18.8	[–3.1, 111.1]
CAD/USD .73 100d	–11.86	–0.8	0.5	1.0	0.8	[–2.5, 1.0]
CAD/USD .74 30d	–11.01	–1.4	1.0	1.6	0.8	[–2.3, 1.6]

Table 6: nVaR analysis for the FX portfolio (1,000 scenarios)

polynomial approximation to the nTRP. This analysis is summarized in Table 8. The results of the polynomial approximation appear in parentheses.

The magnitudes of the nVaR contributions are quite large, ranging from –2387% (Mitsubishi Psep30 800) to 2151% (Komatsu Cjun2 670). This is due to the fact that the portfolio is highly-leveraged and well-hedged, so that the risks incurred by individual positions tend to offset each other to a large extent. In particular, considering the relative sizes of the positions in the portfolio, note that the Komatsu Cjun2 670 position stands to gain considerably if the market appreciates while the Mitsubishi Psep30 800 position acts in the opposite manner. More generally, as indicated by their negative contributions, the two short calls and the two long puts act as a hedge for the portfolio, protecting against drops in the NIKKEI index.

Instrument	Type	Days to Maturity	Strike Price (x 10^3 JPY)	Position (x 10^3)	Value (x 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 7: NIKKEI portfolio

Instrument	nVaR Contribution $nC(x_i)$ (%)	Marginal nVaR $\frac{\partial\, nVaR(x)}{\partial x_i}$ (JPY)	Current Position x_i (x 10^3)	Best Hedge Position x_i^{nbh} (x 10^3)	nVaR Reduction (%)
Komatsu Cjun2 670	2151 (1094)	2727 (1028)	22.5	20.1 (20.0)	41.3 (43.0)
Komatsu Cjun2 760	678 (344)	2576 (970)	7.5	5.0 (4.8)	42.2 (44.2)
Mitsubishi Csep30 836	477 (247)	1699 (653)	8.0	3.9 (4.6)	34.8 (35.3)
Mitsubishi EC 6mo 860	232 (179)	575 (329)	11.5	9.1 (7.8)	22.9 (23.0)
Komatsu	202 (101)	2300 (857)	2.5	–0.4 (–0.1)	35.1 (35.0)
Mitsubishi	149 (75)	2119 (790)	2.0	–1.2 (–0.8)	35.1 (35.0)
Komatsu Paug31 760	53 (28)	–152 (–60)	–10.0	33.2 (28.7)	35.2 (35.7)
Komatsu Cjul29 900	–51 (–26)	52 (20)	–28.0	–150.5 (–121.0)	34.7 (35.2)
Komatsu Paug31 830	–237 (–119)	–675 (–252)	10.0	19.8 (19.0)	34.6 (35.4)
Mitsubishi Cjul29 800	–1166 (–591)	2078 (781)	–16.0	–19.3 (–18.8)	34.8 (35.0)
Mitsubishi Psep30 800	–2387 (–1232)	–1702 (–651)	40.0	44.2 (43.5)	34.7 (35.4)

Table 8: Analysis of the NIKKEI portfolio based on 1,000 scenarios (results of polynomial approximation in parentheses)

The contributions (and the marginal nVaRs) calculated using the polynomial approximation to the nTRP are roughly half the size of those obtained by the pure non-parametric approach. As will be discussed shortly, this is due to the smoothing effects of the approximation. Note, however, that the relative sizes of the contributions among all instruments are the same in both cases (i.e., they yield an identical ranking of the instruments in terms of the nVaR contribution).

Based on the marginal nVaR, the most attractive possibilities for lowering the overall portfolio risk include any one of the following trades: reducing the current holdings in Komatsu Cjun2 670 or Komatsu Cjun2 760, selling one of the common stocks, or shorting additional calls on Mitsubishi (i.e., Mitsubishi Cjul29 800). Purchasing additional units of Mitsubishi Psep30 800 is also a promising option.

If it is feasible to hold an instrument at its best hedge position, then Komatsu Cjun2 760 offers the greatest potential for reducing risk (i.e., a reduction of 42.2%). Note the close agreement between the best hedge positions as determined by the nTRP and by the polynomial approximation.

Smoothing and the polynomial approximation

Figure 8 shows the nTRP and its polynomial approximation for one of the Mitsubishi call options. At the current position (11,500), the nTRP is more steeply sloped than the polynomial, which results in a larger marginal nVaR (i.e., 575 versus 329). Note that smoothing counteracts discrepancies caused by the piecewise linearity of the nTRP. This is particularly evident at a position of 7,000; here, the nTRP slopes upwards, implying a positive marginal nVaR, while the negative marginal nVaR derived from the approximation is more consistent with the general shape of the trade risk profile.

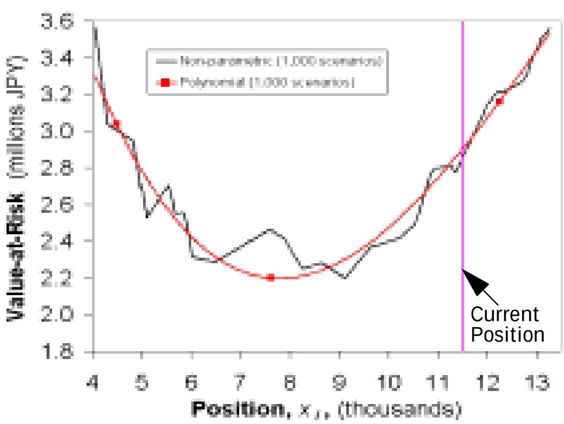

Figure 8: nTRP and polynomial approximation for Mitsubishi EC 6mo 860

To illustrate the effects of increasing the number of scenarios on the nTRP, the NIKKEI portfolio was also simulated over 2,000 and 5,000 scenarios. These scenario sets are obtained by first adding 1,000, and then a further 3,000 scenarios to the initial scenario set. Figure 9 plots the 5,000-scenario nTRP along with each of the polynomial approximations. The curves tend to shift downwards, suggesting a smaller nVaR, as the number of scenarios is increased. However, we note that they remain within the 95% confidence interval for nVaR, calculated for the 1,000-scenario simulation, at the current (11,500) and best hedge positions (9,100). While increasing the number of scenarios creates more segments in the nTRP (compare Figures 8 and 9), sampling error remains a concern even at the 5,000-scenario level (i.e., one can find segments on the nTRP whose slopes are inconsistent with those of the polynomial approximation). In contrast, the approximations tend to provide more consistent gradient information.

It should be noted, however, that even when using the polynomial approximation, we observe inconsistencies in the nVaR contributions (i.e., a position that contributes positively to VaR in one analysis is found to have a negative contribution in another). Since the portfolio nVaR is extremely small relative to the individual position nVaRs in this case, slight errors in the polynomial approximations may combine to change the sign of the denominator in Equation 12. This should only be viewed as a concern for well-hedged, highly-leveraged portfolios, such as the one considered in this example, rather than for portfolios containing options in general.

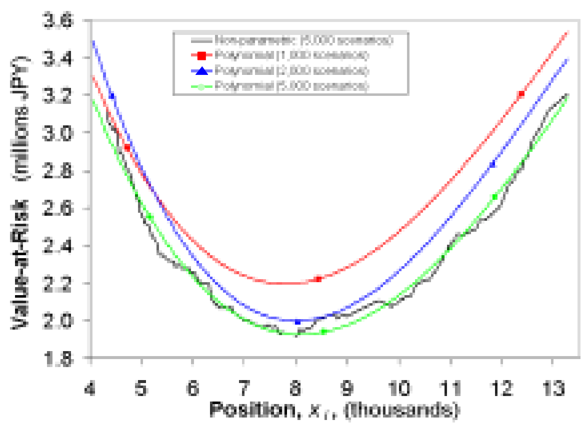

Figure 9: nTRP and polynomial approximations for Mitsubishi EC 6mo 860

Conclusions

This paper has examined tools for VaR-based risk management. Tools for decomposing VaR, assessing its marginal impacts and constructing best hedges, allow managers to understand the sources of risk better and to manipulate the portfolio to effect the desired changes in risk. The analytical techniques that derive from the parametric, or delta-normal, VaR form the basis of a risk manager's toolkit when dealing with portfolios of linear instruments. The simulation-based tools developed in this paper extend these capabilities to portfolios that contain non-linearities or are subject to non-normal market distributions. An attractive feature of these methods is their need for only a single simulation to obtain the mark-to-future values of the instruments. Furthermore, it is straightforward to incorporate new instruments into the analysis by simulating them independently of the portfolio itself. Thus, while our analyses considered trading only those instruments currently held in the portfolio, it extends naturally to encompass so-called incremental VaR. Specifically, it is only necessary to simulate the additional instruments to be considered. Once their mark-to-future values have been obtained, the instruments can be easily incorporated in any marginal nVaR analysis by assigning them a current position of zero.

We note that sampling errors can occasionally yield inconsistent results under the simulation-based approach to VaR. Hence, we propose fitting a smooth curve to the (piecewise-linear) trade risk profile to obtain more robust estimates of the VaR contribution, marginal VaR and best hedge positions. The techniques are demonstrated on a portfolio of European options that is poorly-suited for parametric analysis. The results show that a smooth approximation to the nTRP improves the reliability of marginal VaR estimates, although caution is required when interpreting VaR contributions for well-hedged, highly-leveraged portfolios.

The accuracy of the simulation-based analysis, in light of potential sampling error, is a subject worthy of further investigation. The polynomial approximation to the nTRP is a fairly simple one; the use of splines or other functions may be a preferable approach. In this paper, we have only considered smoothing the trade risk profiles. A more ambitious strategy might seek to fit a probability distribution to the losses themselves and then estimate VaR based on this distribution. This would in fact eliminate the piecewise linearity of the nTRP that is characteristic of the current approach. Furthermore, we anticipate that alternative risk measures, such as expected shortfall, or regret, will exhibit more robust behaviour than nVaR, which relies exclusively on the threshold scenario.

References

Artzner, P., F. Delbaen, J.M. Eber and D. Heath, 1998, "Coherent measures of risk," Working Paper, Institut de Recherche Mathématique Avancée, Université Louis Pasteur et C.N.R.S.

Dowd, K., 1998, "VaR by increments," *Enterprise Wide Risk Management Special Report, Risk*: 31-32.

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

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

Litterman, R., 1996a, "Hot spots and hedges," Risk Management Series, Goldman Sachs.

Litterman, R., 1996b, "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.

Markowitz, H., 1952, "Portfolio selection," *Journal of Finance* 7(1): 77-91.

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

RiskMetrics™ Technical Document, 4th Edition, J.P. Morgan, December 1996.

Sharpe, W., 1964, "Capital asset prices: a theory of market equilibrium under conditions of risk," *Journal of Finance* 19(3): 425-442.

Appendix

Calculating parametric VaR

Consider a portfolio with N holdings that is exposed to W market risk factors. Each instrument in the portfolio is decomposed into a set of risk factor positions so that the change in the instrument's value, Δv_i, can be expressed linearly in terms of the risk factor returns:

$$\Delta v_i = \sum_{k=1}^{W} m_k^i r_k \qquad \text{(A1)}$$

Recall that the vector $\boldsymbol{m}^i$ is the VaR map of instrument i. We can express the change in the value of the portfolio as the sum of the changes in the values of its holdings:

$$\Delta V(\boldsymbol{x}) = \sum_{i=1}^{N} x_i \sum_{k=1}^{W} m_k^i r_k$$

From the definition of the portfolio VaR map (Equation 2), we can write Equation A1 more compactly as

$$\Delta V(\boldsymbol{x}) = \boldsymbol{m}(\boldsymbol{x})^T \boldsymbol{r}$$

Note that $\Delta V(\boldsymbol{x})$ is normally distributed with mean zero and variance $\boldsymbol{m}(\boldsymbol{x})^T \mathbf{Q}^* \boldsymbol{m}(\boldsymbol{x})$, so that the $100(1-\alpha)\%$ VaR is

$$VaR(\boldsymbol{x}) = Z_\alpha \sqrt{\boldsymbol{m}(\boldsymbol{x})^T \mathbf{Q}^* \boldsymbol{m}(\boldsymbol{x})}$$

Defining $\mathbf{Q} = Z_\alpha^2 \mathbf{Q}^*$ (e.g., $Z_{.05} = 1.645$) yields Equation 1.

Parametric trade risk profile

To construct the trade risk profile for instrument i, fix the positions in all instruments other than i to their current values and consolidate them into a base portfolio position x_l. Denote the single-unit volatilities of instrument i and the base portfolio by σ_i and σ_l, respectively, and their correlation by ρ_{il}. The volatility of the portfolio is

$$\sigma(\boldsymbol{x}) = \sqrt{(x_i\sigma_i)^2 + (x_l\sigma_l)^2 + 2\rho_{il}(x_i\sigma_i)(x_l\sigma_l)}$$

Since

$$VaR(\boldsymbol{x}) = Z_\alpha \sigma(\boldsymbol{x})$$

it follows that the trade risk profile is a curve of the form

$$f(x_i) = \sqrt{a x_i^2 + b x_i + c} \qquad \text{(A2)}$$

where $a = (Z_\alpha\sigma_i)^2$, $b = 2Z_\alpha^2 \rho_{il}\sigma_i\sigma_l x_l$ and $c = (Z_\alpha \sigma_l x_l)^2$. Differentiating Equation A2 with respect to x_i yields

$$\frac{d\,f(x_i)}{dx_i} = \frac{2ax_i + b}{f(x_i)}$$

Since $f(x_i)$ is strictly positive, the unique minimum occurs at the best hedge position

$$x_i^* = -\frac{b}{2a} = -\frac{\rho_{il}\sigma_l x_l}{\sigma_i}$$

It is straightforward to show that $f(x_i)$ is symmetric around this point.

Simulation-based trade risk profile

To construct the nTRP for instrument i, let us fix the positions in all instruments other than i to their current values. The loss incurred by the portfolio in scenario j (see Equation 8) can be written

$$L_j(x_i) = \Delta V_{ij} + x_i \Delta v_{ij}$$

where ΔV_{ij} includes the losses due to all instruments other than i. If we plot portfolio losses against the position in instrument i, then each scenario j gives rise to a line $L_j(x_i)$ with slope Δv_{ij} and x-intercept $-\Delta V_{ij} / \Delta v_{ij}$. The piecewise linearity of the nTRP follows from the fact that it is composed of segments from these lines. Specifically, the nTRP consists of the segments defined by the threshold scenario at each position. This is illustrated in Figure A1, which shows a nTRP in which the threshold scenario is always the one with the third-largest loss. Note that each nTRP segment lies on the third line from the top.

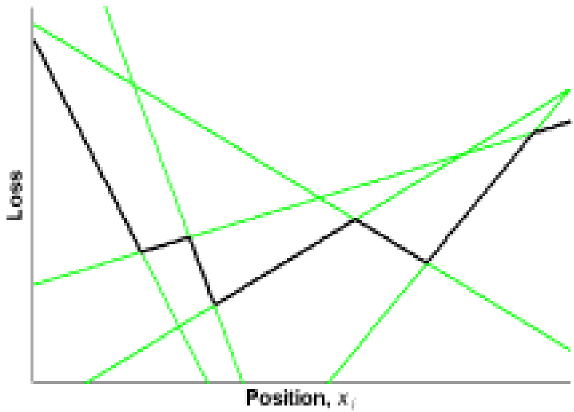

Figure A1: Example of a nTRP (in bold)

In general, the threshold scenario can change whenever the line $L_{s^\circ}(x_i)$ intersects that of another scenario. An algorithm that constructs the trade risk profile and finds the best hedge position is described in Mausser and Rosen (1998).