

Below is a “kind” letter from JOIM asking me to remove “their” copy, and the version of the paper prior to submission.

THANK YOU JOIM! You are too kind.

How would you get people to contribute to your journal their own creation is a mystery to me. Good luck (not really...).

Boudoukh Jacob

From: Christine Proctor [christine@gfong.com] on behalf of JOIM Editor [editor@joim.com]
Sent: 20:21 2005 יום שישי 21 אוקטובר
To: rwhitela@stern.nyu.edu; Boudoukh Jacob
Subject: JOIM

<http://pages.stern.nyu.edu/~jboudouk/maxvarjoim.pdf#search='JOIM%202004'>

As per the JOIM Copyright Agreement: author (s) agrees that any and all content will not be published elsewhere, copied and / or placed on the Internet without written permission from JOIM.

Per my records, JOIM has not given permission to place this article on the above link.

We would appreciate if you would remove this immediately. We would like to suggest placing the article abstract and article reference to the JOIM site (www.joim.com).

Thank you for your cooperation.

Cordially,

Christine Proctor
Editorial Assistant
www.joim.com

MaxVaR:

Long Horizon Value at Risk in a Mark-to-Market Environment¹

Jacob Boudoukh², *Arison School of Business, IDC, Stern School of Business, NYU, and NBER*

Matthew Richardson, *Stern School of Business, NYU, and NBER*

Richard Stanton, *Haas School of Business, UC Berkeley*

Robert F. Whitelaw, *Stern School of Business, NYU, and NBER*

Abstract:

The standard VaR approach considers only terminal risk, completely ignoring the sample path of portfolio values. In reality interim risk may be critical in a mark-to-market environment. Sharp declines in value may generate margin calls and affect trading strategies. In this paper we introduce the notion of ***MaxVaR***, analogous to VaR in every way except it quantifies the probability of seeing a given loss on *or before* the terminal date rather than at the terminal date. Under standard set of assumptions we provide a simple formula for MaxVaR and examine the ratio of MaxVaR to VaR. For reasonable parameterizations *MaxVaR* may exceed VaR by over 40%. MaxVaR exceeds VaR by as much as 80% or more for high Sharpe Ratio hedge-fund-like sets of portfolio return distribution.

¹ An Excel spreadsheet with a calculator for the implicit solution for MaxVaR is available upon request. We thank an anonymous referee.

² Contact author. Email: jboudouk@stern.nyu.edu.

1. The Problem with Long Horizon VaR

In spite of the popularity of value at risk, there are many unresolved issues from both a methodological and an implementation perspective. Methodological issues revolve mainly around interpretation. For example, VaR tells us how frequently we expect to see a loss greater than a certain threshold value. VaR, however, says little about the distribution of such losses when they do occur.³ With respect to implementation, there are difficulties in estimating VaR due to the non-normal behavior of asset returns. Fat tails, skewness and correlation breakdown, typical properties of asset returns, result in VaR calculations under the normality assumption often yielding estimates that err on the conservative side. Nonparametric models such as historical simulation, a theoretically suitable remedy, often fail due to sampling errors.^{4 5}

The increasingly popular use of VaR in long-horizon (or multiperiod) settings is even more difficult to interpret. In addition to many concerns that are valid for short horizons, when we use VaR for long horizons we need to add a particularly unpalatable assumption that there is no intermediate trading between the time the VaR is calculated and the terminal measurement horizon. While the absence of intermediate trading is a strong assumption, there are many reasons to adopt it, not least because there is no clear alternative. In reality when trading results turn sour a trader or an institution may choose to or may have to unwind positions. On the other hand, it may also “go for broke” and increase positions. Either change may, however, be difficult to execute due to liquidity constraints. Hence, the assumption of a fixed position may be thought of as reasonable for particularly large positions. In this paper we make the simplifying assumption of no

³ See “Expect the Worst -- Rethinking the Value at Risk Concept using Worst Case Scenario Analysis and its Implications for Risk Management,” Jacob Boudoukh, Matthew Richardson, and Robert F. Whitelaw, *Risk*, September 1995.

⁴ For a discussion of some of these difficulties see “The Best of Both Worlds: A Hybrid Approach to Calculating Value at Risk,” Jacob Boudoukh, Matthew Richardson and Robert F. Whitelaw, *Risk*, May 1998.

⁵ For a general discussion of VaR see “Understanding Market, Credit, and Operational Risk – the Value at Risk Approach”, Linda Allen, Jacob Boudoukh and Anthony Saunders, 2004, Blackwell Publishing.

change in portfolio composition throughout. This is a useful assumption also in order to preserve comparability with the standard VaR concept.

In this article we address a specific methodological problem that arises when VaR is used for long-horizon risk measurement and management purposes. The problem is that VaR calculations focus on the distribution of a portfolio's value only at the terminal date, ignoring the path of portfolio value prior to that date. The key question we ask is what happened in the interim? In a mark-to-market environment this is a very relevant question, particularly due to the increased use of long-horizon VaR in the money management and insurance industries, as well as by regulators and investment banks. As opposed to VaR, the loss we may see at a given horizon with a certain probability, we call the loss we may see *on or before* the given horizon MaxVaR. At the 5 percent level and for typical equity market parameters MaxVaR exceeds VaR by over 40%. For higher Sharpe Ratio strategies MaxVaR may exceed VaR by over 80%.


2. Framework

Commonly used VaR calculations focus on the distribution of asset returns at a given horizon. For example, regulators may require institutions to report their 10-day VaR, or a fund manager may want to know the quarterly VaR of a portfolio. Statements about VaR are, of course, probabilistic, and provide information about the tail of the distribution at the given horizon. However, it is also valuable to the decision maker, viewing a daily mark-to-market of his/her portfolio, to know information about the distribution of asset returns on or before that horizon. While one could argue that the horizon is arbitrary in any case, the VaR at any horizon is somewhat deceiving when portfolio values are monitored day-to-day.

VaR is especially misleading if a margin call can occur given a certain mark-to-market value of the portfolio. In such a case the use of VaR vis-à-vis the capital of the firm/hedge fund/trading strategy is completely misguided. We are, in fact, interested in the probability of going below a certain value *on or prior to* the chosen horizon.

To implement this notion we define MaxVaR as the loss in portfolio value that will be exceeded with a given probability level on *or before* the given horizon. MaxVaR is, therefore, immediately comparable to the standard VaR that focuses on terminal values only. The ratio that we calculate (MaxVaR divided by VaR) is simply an adjustment factor to the standard VaR, which is readily usable as an extension of the standard calculation. We therefore use all the standard distributional assumptions here.

Let the value of the asset (or the portfolio, for that matter), S_t , follow a log-normal diffusion process with annualized expected return μ and volatility σ . The value at the end of period T , S_T , is the variable of interest when calculating the T period VaR. In order to quantify MaxVaR, define the minimal sample-path value of S_t between 0 and T to be $\underline{S}_{(0,T]}$. That is, $\underline{S}_{(0,T]} \equiv \min\{S_t, t \text{ in } (0,T] \}$. We focus on the distribution of $\underline{S}_{(0,T]}$. (These quantities are illustrated graphically in the Figure 1.) Similarly we can consider rates of return. Let $R_t = \ln(S_t/S_0)$ be the (continuously compounded) return from 0 to t . Let $\underline{R}_{(0,T]}$ be defined as $\underline{R}_{(0,T]} = \ln(\underline{S}_{(0,T]}/S_0)$, i.e., the minimum cumulative return over the period.



Further, define $m = \mu - \sigma^2/2$, i.e., the annual expected continuously compounded return. For a given tail value z , the probability α_{VaR} that $R_T \leq z$ is $\Phi\left(\frac{z - mT}{\sigma\sqrt{T}}\right)$, where $\Phi(\cdot)$ is the cumulative standard normal.

Theorem:

$$\alpha_{MaxVaR} \equiv Prob[\underline{R}_{(0,T]} < z] = \Phi\left(\frac{z - mT}{\sigma\sqrt{T}}\right) + e^{2mz/\sigma^2} \Phi\left(\frac{z + mT}{\sigma\sqrt{T}}\right) \quad (for\ z < 0)$$

Proof:

The result follows from the properties of the first passage time of a Brownian motion with drift⁶.

To get some intuition for this result consider the simpler case where $m=0$, i.e., there is no expected drift in the (log) price process. In this case, the above formula simplifies to

$$\alpha_{MaxVaR} = Prob[\underline{R}_{(0,T]} < z] = 2\Phi\left(\frac{z}{\sigma\sqrt{T}}\right) = 2\alpha_{VaR} \quad (for\ z < 0).$$

The MaxVaR probability of hitting z or lower on or before T is, interestingly, exactly twice the probability associated with hitting z at the horizon, T . i.e., the VaR. In this simpler case the formula for the MaxVaR is based on the well-known “reflection principle” -- for every process with a sample path that touches the value z and declines further there is one that rises from that z point on. These could be thought of as mirror image processes. To summarize, the z percent VaR is the $2z$ percent MaxVaR. Of course, this does not imply that the $z\%$ MaxVaR is twice the $z\%$ VaR.

⁶ See, for example, Ioannis Karatzas and Steven E. Shreve, *Brownian Motion and Stochastic Calculus*, second edition, Springer (New York), Graduate Texts in Mathematics #113, 1991, pp. 196 – 197. .

3. Results

In Table 1 below we provide a comparison between the standard VaR and the MaxVaR in order to get some intuition for the magnitude of the adjustment which is required in order to account for interim risk. We set $m=0$, i.e., the expected return on the portfolio is assumed to be 0. In this case, while the probability of seeing a return of -1.645 standard deviations or worse at the end of period is 5%, there is a 10% probability of seeing this size move along the sample path prior to the terminal date. To put it differently, there is a 5% probability of seeing an end-of period return of less than -1.645 standard deviations (VaR), and a 5% probability of seeing a return of less than -1.960 standard deviations *or before* the end of the period (MaxVaR). At the 5% level, the ratio of the VaR to the MaxVaR is, therefore, $1.960/1.645=1.192$. That is, the standard VaR calculation can be easily adjusted to account for the fact that value is observed continuously in the interim simply by inflating VaR by 19%. This “inflation” factor, the ratio of MaxVaR to VaR, does not depend on volatility or the horizon, but it declines as the tail probability declines. Thus, the adjustment grows smaller in percentage terms as the tail event becomes less likely.

Table 1: MaxVaR vs. VaR, $m=0$

PROB	VaR	MaxVaR	<u>MaxVaR</u> VaR	MaxVaR (N=10)
5%	$1.645 \sigma\sqrt{T}$	$1.960 \sigma\sqrt{T}$	1.192	$1.802 \sigma\sqrt{T}$
2.5%	$1.960 \sigma\sqrt{T}$	$2.241 \sigma\sqrt{T}$	1.144	$2.090 \sigma\sqrt{T}$
1%	$2.326 \sigma\sqrt{T}$	$2.576 \sigma\sqrt{T}$	1.107	$2.420 \sigma\sqrt{T}$

In the last column of the table we also provide an example of this calculation when the interim sampling is discrete. Suppose, for example, the process is sampled ten times (N=10) over the relevant horizon. The discrete MaxVaR will be lower than the continuous

sampling MaxVaR since we might “miss” the lowest point.⁷ We use 50,000 simulations to determine the appropriate value. The case of $N=10$ is of particular interest due to the fact that regulators often look at a 10-day VaR as a measure of banks’ trading risk (see, e.g., the 1998 Basel capital requirements). These market risk-related capital requirements consider the two-week (i.e., ten day) VaR as the basis, but, we argue, ignore the mark-to-market of the trading portfolio in the interim. The appropriate adjustment for daily mark-to-market appears in the table above. With discrete sampling at this frequency, the necessary percentage adjustment is slightly less than half of that needed for continuous sampling.

Assuming the drift is zero is unrealistic for most practical examples, although it is still useful for two reasons. First, $m=0$ provides a lower bound on the necessary adjustment factor for positive drift processes. As the drift increases, the ratio of MaxVaR to VaR also increases because interim returns have lower means than their terminal horizon counterparts. Second, $m=0$ provides the limit of the adjustment as the horizon shrinks -- for very short horizons the drift is relatively unimportant.

To get an idea of the potential effect of a positive drift for longer horizons, Table 2 reports results for $\sigma=15\%$, $T=1$, and $\mu=10\%$ or 15% . To make comparison easier, the values are reported in units of standard deviation ($\sigma\sqrt{T}$), as in Table 1. However, in contrast to Table 1, the results are not independent of σ and T (we will expand on this issue below).

There are a couple of results worth noting. First, and most obvious, both the VaR and MaxVaR are reduced relative to the levels reported in Table 1 due to the positive drift. For example the 5% MaxVaRs are 1.493 and 1.262 standard deviations for μ equal to 10% and 15%, respectively, relative to 1.960 in the no drift case. Second, and more important, the MaxVaR/VaR ratio can increase dramatically as the drift increases. For example, at the 5% level, the MaxVaR is now more than 41% or 75% greater than the corresponding VaR

⁷ A similar issue arises in the case of barrier options, where the question of whether the knock in/out is a function of the end of the trading day closing price or at any point during the trading day turns out to make a difference for pricing these derivatives.

for the two drift values—a very economically significant adjustment factor. This increase is due to the fact that the positive drift has less influence on returns over intermediate horizons. Thus MaxVaR is reduced less than VaR, and the ratio increases. This result demonstrates the fact that the use of VaR instead of MaxVaR as a risk measure is especially problematic for high Sharpe Ratio portfolios, e.g., hedge funds.

Table 2: MaxVaR vs. VaR, $\sigma=15\%$, $T=1$

PROB	μ	VaR	MaxVaR	<u>MaxVaR</u> VaR
5%	10%	$1.053 \sigma\sqrt{T}$	$1.493 \sigma\sqrt{T}$	1.417
2.5%	10%	$1.368 \sigma\sqrt{T}$	$1.752 \sigma\sqrt{T}$	1.281
1%	10%	$1.735 \sigma\sqrt{T}$	$2.067 \sigma\sqrt{T}$	1.191
5%	15%	$0.720 \sigma\sqrt{T}$	$1.262 \sigma\sqrt{T}$	1.753
2.5%	15%	$1.035 \sigma\sqrt{T}$	$1.504 \sigma\sqrt{T}$	1.453
1%	15%	$1.401 \sigma\sqrt{T}$	$1.801 \sigma\sqrt{T}$	1.285

It is important to note that the results in Table 2 depend on all three parameters— μ , σ , and T . In fact, the key variable is the normalized drift over the relevant horizon, i.e., $\frac{(\mu - \sigma^2/2)T}{\sigma\sqrt{T}} = \frac{(\mu - \sigma^2/2)}{\sigma}\sqrt{T}$. As this quantity increases, so does the MaxVaR adjustment factor. Therefore, the effect of considering interim risk is increasing in μ and T and decreasing in σ .

4. Conclusion

We provide an important extension to the well-known measure of risk, VaR, to account for the fact that the sample path of portfolio values is observable via mark-to-market, and is of interest because it may affect trading activity in the portfolio (e.g., cause liquidation or

stop-loss triggers). The concept of MaxVaR is easy to quantify and we provide a ratio of the standard VaR to MaxVaR for a number of illustrative cases.

Further extensions to this concept may be of interest. In particular, our derivation uses the standard assumption that returns are independent log-normals with constant parameters. Various violations of these assumptions may change the MaxVaR calculation. First, asset returns are known to be fat tailed. It is common for risk management systems to ignore fat tails and assume that long horizon returns are normal. From a practical perspective this is a plausible assumption, due to the fact that fat tails are likely to “wash out” by averaging over long horizons (this is simply a result of the law of large numbers under some mild assumptions). This may not be such a good assumption for interim values. While the effect of unusually large moves may fade out at long horizons, they may still be important at short horizons and lead to increases in MaxVaR.

Second, if asset returns are positively or negatively serially correlated, i.e., asset prices experience return continuations or reversals, then much in the same way that long horizon VaRs need to be calculated carefully, taking these time series properties into account will also affect the MaxVaR.

Finally, in the presence of stochastic volatility similar issues may arise, where periods of high volatility may generate extreme moves that, if we ignore them, we may understate the MaxVaR. These topics are left for future research.