How valuable is your VaR? Large sample confidence intervals for normal VaR

Received (in revised form): 14th December, 2010

Franck Moraux
is a Professor of Finance at the University of Rennes 1 Graduate School of Management. His research interests cover financial markets, risk management, quantitative finance, option theory and applications. He has published various research articles and a book in these fields. He is co-editor-in-chief of 'Finance' the official publication of the French Finance Association.

University of Rennes 1, Graduate School of Management, 11 rue Jean Macé, 35000 Rennes, France
Tel: +33 (0)223 237 808; E-mail: franck.moraux@univ-rennes1.fr

Abstract Little is known about the distribution of the 'value-at-risk' (*VaR*) estimate and the associated estimation risk. In the case of the normal *VaR*, the key problem comes from the fact that it is estimated using a couple of parameters whose estimates are distributed differently. Previous research has either neglected uncertainty around the mean parameter, or resorted to simulations. By contrast, this paper derives analytical results for the normal *VaR* with the help of asymptotic theory and the so-called 'delta method'. Properties of the estimation errors are then explored in detail and the *VaR* estimation risk is broken down into its various components. It is then shown, among other things, that the fraction of error owing to mean uncertainty is limited in a prudential context. In other words, the approximate approach defended by Jorion and Chappell and Dowd is shown to still be relevant.

Keywords: *value-at-risk, estimation risk, confidence interval, large sample*

INTRODUCTION

The 'value-at-risk' (*VaR* hereafter) is still nowadays an important risk measure for risk management purposes even if significant limitations and warnings have been exposed in the recent literature (see, among many others, Artzner *et al.*[1]). Among existing methods, the parametric approach is rather popular for computing *VaRs*. Among parametric specifications, the normal *VaR* is often used (at least) as a benchmark. (The assumption of normal returns is not discussed in the rest of the paper. Nevertheless, if portfolio or asset returns are not normally distributed, the use of standard normal *VaRs* for risk management purposes should be questioned. Typically, Christoffersen and Gonçalves[2] consider portfolio returns with non-constant (conditional) variance. And they show that a re-sampling technique helps to assess the parameter estimation error and construct VaR confidence intervals.)

Surprisingly, even in the standard Gaussian world, little is known about the distribution of the *VaR* estimate and the associated estimation risk. In the normal case, the key problem comes from the fact that the normal *VaR* involves a couple of parameters (a mean and a standard deviation) whose estimates are distributed differently. It is well known in statistics that the sample mean is exactly normally distributed whereas the sample variance is exactly chi-square distributed. According to Kendall and Stuart[3] (quoted by Chappell and Dowd[4]), finding the exact confidence interval for a function of two estimates can bear 'very considerable difficulty'. This can partly explain why, in past research, authors have either neglected the uncertainty related to the mean parameter or have resorted directly to simulations.

Typically, Jorion[5,6] investigates the estimation error of normal *sigma*-based *VaR* where the mean parameter is set to zero. Chappell and Dowd[4] provide the exact confidence interval for normal *VaR* when the mean is supposed to be known without uncertainty. These approaches provide practical and rather simple ways to compute confidence intervals. Yet, the assumption they make can be viewed as rather strong. To relax it, Dowd[7] proposes a pragmatic approach relying on simulation (see Cotter and Dowd[8] for further developments in that direction). Summing up, there is no analytical result for the exact distribution of the normal VaR estimate, nor analytical expression for associated confidence intervals.

Interestingly, one can observe that all these approaches refer to exact statistical results such as the exact sampling distributions. By contrast, the present paper derives various analytical approximations by exploiting the asymptotic theory of statistics and the so-called 'delta method'. The asymptotic theory tells us that every sample estimate approaches normality in the limit as the sample size grows, while the delta method allows us to describe the variance of the normal *VaR*. Because this paper accounts for the 'mean' uncertainty, it can be viewed as an extension of the analytical approaches exposed in Jorion[5,6] and Chappell and Dowd.[4] Avoiding simulations, it offers a pragmatic alternative to Dowd.[7]

The rest of the paper proceeds as follows. The second section presents the framework and formalises the motivations. The third section derives the key results on the asymptotic behaviour of the *VaR* estimates and provides a couple of large sample confidence intervals for normal *VaR*. The next section discusses operational and managerial implications of these results. The final section aims at evaluating the large sample results by comparison to Chappell and Dowd[4] in a framework where the mean uncertainty is neglected.

NOTATIONS, FRAMEWORK AND MOTIVATIONS

One can denote by P_s and $P\&L_t(\Delta t) = P_{t+\Delta t} - P_t$ values of the portfolio and its associated profit and loss over the next investment period $[t, t+\Delta t]$. At time t, profits and losses as well as 'ex ante' portfolio returns (computed by $P\&L_t(\Delta t)/P_t$) are random by nature. The (return) *VaR* is then defined by the strictly positive real number such that:

$$\Pr\left(\frac{P\&L_t(\Delta t)}{P_t} \leq -VaR_\alpha\right) = 1 - \alpha$$

where α is the *VaR* confidence level. In applications, α lies between 95 per cent and 100 per cent (this latter value being excluded). It is typically equal to 95, 97.5, 99 or 99.9 per cent in 'real life' risk management contexts. The normal *VaR* asserts that returns are normally distributed. Denoting by μ the 'true' expected return and σ its 'true' standard deviation, the normal *VaR* is then given by:

$$VaR_\alpha = -\mu - \sigma q_\alpha = f(\mu, \sigma) \qquad (1)$$

where q_α stands for the $(1-\alpha)-$ quantile of the standard normal distribution. Denoting by N the cumulative probability function of the Gaussian distribution, $q_\alpha = N^{-1}[1-\alpha]$. The 'true' normal *VaR* is therefore a bilinear function of the 'true' (unknown) mean and the 'true' (unknown) standard deviation. It must be stressed that the use of the normal VaR is fully justified when returns are normally distributed. Otherwise, it is just a proxy for the quantile associated to the 'true' data generating process. In this view, equation (1) appears especially credible if the 'true' distribution is well described by a couple of parameters (the mean and the standard deviation) and if it resembles the Gaussian distribution in terms of symmetry and tails. Otherwise, if (eg) the 'true' distribution is asymmetric or fat-tailed, the normal approach may be of limited interest. Whatever the case, even in a normal world, there exists a sampling problem which is caused by the use of observed data. A normal *VaR* estimate is actually computed by:

$$\widehat{VaR}_\alpha = -\widehat{\mu} - \widehat{\sigma} q_\alpha = f(\widehat{\mu}, \widehat{\sigma}) \qquad (2)$$

where $\widehat{\mu}$ and $\widehat{\sigma}$ are (known) sample estimates of the mean and the standard deviation. Note that hereafter the subscript in *VaR* is omitted.

The aim of this paper is to provide the *asymptotic* distribution of the normal *VaR* estimate, given that both estimates ($\widehat{\mu}$ and $\widehat{\sigma}$) carry some degree of uncertainty about 'true' values. This approach is in sharp contrast with previous research. For example, Jorion[5,6] provides the standard error of the *sigma*-based *VaR* estimate, that assumes $\mu = 0$ with certainty. Furthermore, the existing literature dealing with normal *VaRs* has mainly concentrated its attention on *exact* results of statistical theory. Under the normal *iid* assumption, unbiased sample estimates of μ and σ^2 on T observations are known to be *independent* of each other and admit the following exact distributions:

$$\widehat{\mu} \sim N\left(\mu, \frac{\sigma^2}{T}\right), \qquad (3)$$

$$(T-1)\frac{\widehat{\sigma^2}}{\sigma^2} \sim \chi^2_{T-1} \qquad (3')$$

where N denotes the normal distribution and χ^2_{T-1} is the chi-square distribution with $T-1$ degrees of freedom. Hence, as noted by Chappell and Dowd,[4] the chi-square distribution alone provides the exact confidence intervals for *VaR*, when the mean μ is known (or, more precisely, the uncertainty of its estimate is neglected). To relax this latter assumption, Dowd[7] rearranges expressions (3 and $3'$) and states that the normal *VaR* distribution is described by:

$$-N\left(\widehat{\mu}, \frac{(T-1)}{T\chi^2_{T-1}}\widehat{\sigma^2}\right) - q_\alpha\sqrt{(T-1)\frac{\widehat{\sigma^2}}{\chi^2_{T-1}}}.$$

One must however resort to Monte Carlo simulation in order to estimate associated quantiles. Expressions (3) and ($3'$) suggest that exact formulae for the mean and the variance of the normal *VaR* estimate are simple to derive. Yet this is not as easy as it may seem. A reason for this is that the square root of the sample variance is a downward biased estimate of the standard deviation. Consult the Appendix for further details on that point.

LARGE SAMPLE RESULTS FOR NORMAL VAR

The following results benefit from the asymptotic theory of statistics. They are based on a) the asymptotic behaviour of the estimates (of μ and σ^2) and b) the delta method. One solves the confidence interval problem of *VaR* for the 'large sample' case. The next section will investigate theoretical results from a managerial perspective.

Proposition 1. *If portfolio returns are iid and normally distributed, the asymptotic distribution of the (centred and scaled) VaR estimate is given by:*

$$\sqrt{T}\left(\widehat{VaR} - VaR\right) \overset{a}{\sim} N\left(0, V_g\right) \qquad (4)$$

where the 'asymptotic' variance is $V_g = \sigma^2\left(1 + 1/2q_\alpha^2\right)$ *with* α *as the confidence level.* (Note that the asymptotic variance is denoted by V_g (in lieu of a more natural V_a) in order to avoid visual confusion between a and α.)

Proof. Under the assumption that returns are identically and independently normally distributed, the asymptotic distribution of the sample mean is $\sqrt{T}(\widehat{\mu} - \mu) \overset{a}{\sim} N(0, \sigma^2)$ and that of (centred and scaled) sample variance $\sqrt{T}\left(\widehat{\sigma^2} - \sigma^2\right)$ is $N(0, 2\sigma^4)$. Because the normal *VaR* may be rewritten $VaR_\alpha = -\mu - \sqrt{\sigma^2}q_\alpha = g(\mu, \sigma^2)$ the risk measure can be viewed as a (non-linear) function of the variance. *VaR* estimates can then be computed with $\widehat{VaR}_\alpha = g\left(\widehat{\mu}, \widehat{\sigma^2}\right)$ where $\widehat{\mu}$ and $\widehat{\sigma^2}$ are estimates of μ and σ^2. Denoting by θ the column vector $(\mu, \sigma^2)'$ (where $'$ means transpose), the joint asymptotic approximate distribution of the corresponding estimator $\widehat{\theta}$ is given by:

$$\sqrt{T}\left(\widehat{\theta} - \theta\right) \overset{a}{\sim} N(0, V_\theta)$$

with $V_\theta = \begin{pmatrix} \sigma^2 & 0 \\ 0 & 2\sigma^4 \end{pmatrix}$ (remember the independence result between $\widehat{\mu}$ and $\widehat{\sigma^2}$). Viewed as a function of $\widehat{\theta}$, the asymptotic distribution of the $\widehat{VaR} = g\left(\widehat{\theta}\right)$ benefits from the delta method. The delta method states that:

$$\sqrt{T}\left(\widehat{VaR} - VaR\right) \overset{a}{\sim} N\left(0, V_g\right)$$

where:

$$V_g = \left(\frac{\partial g}{\partial \theta}\right)' V_\theta \frac{\partial g}{\partial \theta}$$

with:

$$\left(\frac{\partial g}{\partial \theta}\right)' = \begin{pmatrix} -1 & -\dfrac{q_\alpha}{2\sqrt{\sigma^2}} \end{pmatrix}.$$

This yields to:

$$V_g = \sigma^2 + \frac{2\sigma^4}{4\sigma^2} q_\alpha^2 = \sigma^2\left(1 + \frac{1}{2}q_\alpha^2\right).$$

It is interesting to note that $\sigma^2/2$ in V_g admits a direct interpretation. It is the asymptotic variance of the (scaled) sample standard deviation (var$(\sqrt{T}\widehat{\sigma})$). So, V_g is the sum of two asymptotic terms: the variance of the (scaled) sample mean (var $(\sqrt{T}\widehat{\mu})$) plus the asymptotic variance of the (scaled) sample standard deviation multiplied by the squared normal quantile $\left(q_\alpha^2\right)$.

Various implications may be derived from proposition 1. Standard errors and confidence intervals for large samples are exposed in the two following corollaries. The first corollary focuses on the normal *VaR* estimate (assuming that the 'true' parameters are known). The second one deduces from proposition 1 lower and upper bounds for the (unknown) theoretical or 'true' normal *VaR* when only estimates are known. These results are clearly very useful for controlling and testing normal *VaR*.

Corollary 2. *For a large sample with size T, the (theoretical) standard error of the* $\widehat{VaR}$ *may be written:*

$$S.E.\left(\widehat{VaR}\right) = \frac{\sigma}{\sqrt{T}}\sqrt{1+\frac{1}{2}q_\alpha^2} = |VaR+\mu|\frac{\sqrt{2+q_\alpha^2}}{|q_\alpha|\sqrt{2T}} \tag{5}$$

and the analytical (theoretical) 'large sample' confidence interval for the VaR estimate is:

$$VaR + S.E.\left(\widehat{VaR}\right)q_{\frac{1-\beta}{2}} < \widehat{VaR} < VaR + S.E.\left(\widehat{VaR}\right)q_{\frac{1+\beta}{2}} \tag{6}$$

with β *being the chosen confidence level for the VaR estimate.*

Corollary 3. *For a large sample with size T, the 'large sample' confidence interval for the normal VaR is:*

$$\widehat{VaR} - \overline{S.E.\left(\widehat{VaR}\right)}q_{\frac{1+\beta}{2}} < VaR < \widehat{VaR} - \overline{S.E.\left(\widehat{VaR}\right)}q_{\frac{1-\beta}{2}} \tag{7}$$

where the estimated standard error of the $\widehat{VaR}$ *is given by:*

$$\overline{S.E.\left(\widehat{VaR}\right)} \approx \sqrt{\widehat{\sigma}^2}\frac{\sqrt{2+q_\alpha^2}}{\sqrt{2T}} = \left|\widehat{VaR}+\widehat{\mu}\right|\frac{\sqrt{2+q_\alpha^2}}{|q_\alpha|\sqrt{2T}}. \tag{8}$$

with β *being the chosen confidence level for the VaR estimate.* (The confidence level β associated to the *VaR* estimate must not be confused with α which is the confidence level in the risk measure itself. For a one-side test, the single bound is easy to compute.)

DISCUSSIONS AND MANAGERIAL IMPLICATIONS

This section discusses the results of the preceding section with real-life concerns in view. First of all, standard errors of the normal *VaR* estimate, exposed in (5) and (8), appear inversely proportional to the square root of the sample size $(\sqrt{T})$. This means that a financial risk manager who wants to improve by a factor of two the estimation accuracy needs to find four times as much data *ceteris paribus*. Standard errors also increase with the confidence level α and grow proportionally with the relative *VaR* $(VaR + \mu)$. Equations (5) and (8) assess quantitatively how the accuracy decreases, or how the estimated risk measure is riskier, as one stands further in the tail. To illustrate this, Table 1 reports values of $\sqrt{2+q_\alpha^2}/|q_\alpha|\sqrt{2T}$ for different values of α and T. Standard errors of normal *VaR* can then be computed by multiplying the corresponding relative *VaR* estimates to values reported in Table 1, by using (8). For instance, for a $N(0,1)$ distributed return and $T = 250$, standard errors of the normal *VaR* are $1.64485 \times 0.059 \approx 0.0970$ if $\alpha = 95\%$ and $2.32635 \times 0.052 \approx 0.1210$ if $\alpha = 99\%$. One can verify in Table 1 that values decrease by a factor of 2 when the sample size T increases by a factor of 4. When 24 observations are used, standard errors represent at least 15.9 per cent of the relative *VaR*. It is lower than 5 per cent for 250 observations.

The '2' term under the square root of equation (5) is directly related to the mean uncertainty. This term therefore disappears if μ is known with certainty.

Table 1: Asymptotic standard errors of normal value-at-risk for combinations of confidence level and sample size (expressed in percentage of the relative value-at-risk)

α	T 12	24	48	52	250	500	1000
90.0%	0.304	0.215	0.152	0.146	0.067	0.047	0.033
92.5%	0.286	0.202	0.143	0.137	0.063	0.044	0.031
95.0%	0.269	0.190	0.135	0.129	0.059	0.042	0.029
97.5%	0.252	0.178	0.126	0.121	0.055	0.039	0.028
99.0%	0.239	0.169	0.119	0.115	0.052	0.037	0.026
99.9%	0.224	0.159	0.112	0.108	0.049	0.035	0.025

Note: The table displays different values of $\sqrt{2+q_\alpha^2}/|q_\alpha|\sqrt{2T}$ useful to compute asymptotic standard errors of normal value-at-risk.

The standard error then becomes $\sqrt{\widehat{\sigma^2}}/\sqrt{2T}|q_\alpha|$ which is the result of Jorion[5,6] (known to be $\widehat{\sigma}/\sqrt{2T}|q_\alpha|$). The uncertainty about the mean affects the overall asymptotic standard error. Inspired from Jorion,[5,6] Figure 1 compares standard errors of the estimated *VaR* for a one-year sample of daily data ($T = 250$) if one neglects the mean uncertainty ('J' line for Jorion's case) or not. As expected, total standard errors are systematically greater. Interestingly, the difference is analytically given by $\widehat{\sigma}/\sqrt{2T}\left(\sqrt{2+q_\alpha^2}-|q_\alpha|\right)$ whereas the ratio $((\sqrt{2+q_\alpha^2}/|q_\alpha|))$ remains first constant with respect to the volatility and secondly greater than 1. Overall, this suggests that, for some confidence levels α, the uncertainty on the estimation of μ should not be underestimated.

The total variance exposed in proposition 1, $V_g = \sigma^2(1+1/2q_\alpha^2)$, may furthermore be decomposed in two terms: σ^2 and $\sigma^2 q_\alpha^2/2$ that correspond to the two sources of uncertainty. As a result, the fraction of error caused by the mean may be assessed by $\sigma^2/V_g = 2/(2+q_\alpha^2)$ whereas the one associated to the standard deviation is given by $q_\alpha^2/(2+q_\alpha^2)$. Interestingly, both fractions are independent of the variance of the underlying return (σ^2). Moreover, they are equal when $2 = q_\alpha^2$, ie for

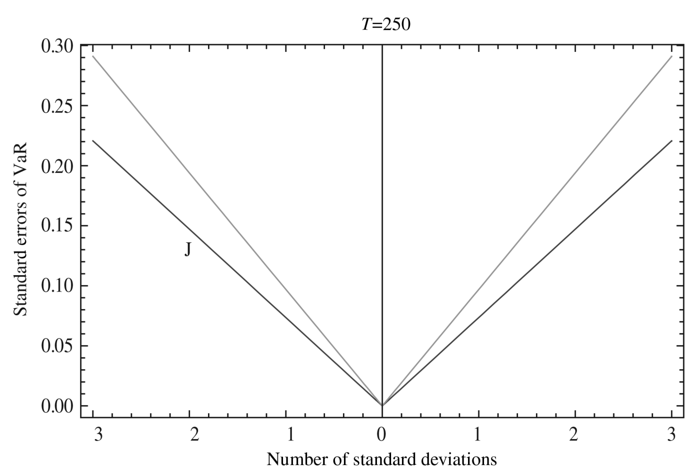

Figure 1: Standard errors for normal value-at-risk
Note: This graph plots standard errors of the normal *VaR* if one neglects the mean uncertainty as in Jorion (the 'J' line, or not).

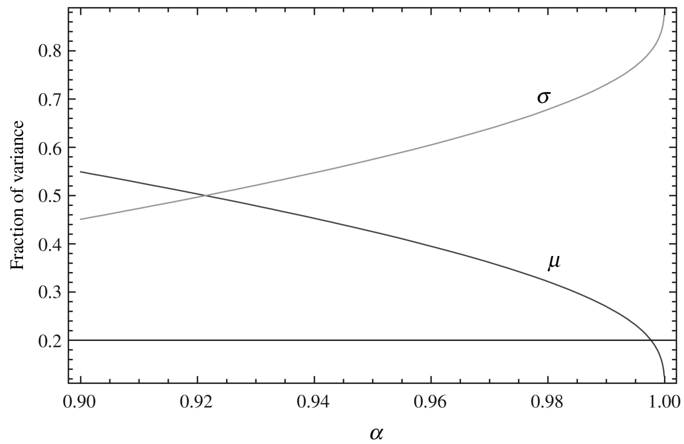

Figure 2: Variance decomposition of the normal value-at-risk estimate
Notes: This figure plots the decomposition of the asymptotic variance of the (centred and scaled) normal value-at-risk as defined in equation (4). The 'σ' line stands for the fraction of variance due to the volatility uncertainty. The 'μ' line corresponds to that of the mean uncertainty. α is the confidence level associated to the 'true' value-at-risk.

$\sqrt{2+q_\alpha^2}/|q_\alpha|\sqrt{2T}$ For $\alpha < \alpha^*$, the overall error is mainly caused by the estimation error of the mean. As a result, in a prudential context (where α is typically above 95 per cent), the estimation error on the mean is not the main source of uncertainty. Figure 2 plots respective contributions of the μ and σ estimation uncertainties on the global variance V_g for different values of α. As expected, lines cross at $\alpha^* \approx 0.92$ where both contributions are equal to 50 per cent. Figure 2 shows that, as α gets larger, the uncertainty on the estimation of σ becomes more significant (relative to that of μ). This suggests that the approximate approach of Jorion[5,6] and Chappell and Dowd,[4] which neglects the mean uncertainty, is especially relevant for the largest values of α.

ASSESSING LARGE SAMPLE RESULTS

To assess the above asymptotic results and have a look at the small sample bias, large sample confidence intervals may be compared to confidence intervals resulting from exact sampling distribution. Clearly, results of Chappell and Dowd[4] appear especially relevant for this exercise even if they focus on the volatility parameter. Actually, the volatility estimate is not normally distributed so that it can be the source of mis-specification. On the contrary, the mean estimate is normally distributed and its uncertainty should not dramatically affect the overall performance of the asymptotic results.

Assuming that the mean is known with certainty ($\Pr[\widehat{\mu} = \mu] = 1$) and exploiting equation ($3'$), Chappell and Dowd[4] state that the $\beta-$confidence interval for the normal VaR_α is:

$$\begin{aligned} -\widehat{\mu} - \sqrt{\widehat{\sigma^2}}\sqrt{(T-1)/q^{\chi^2_{T-1}}_{\frac{1+\beta}{2}}}\,q_\alpha &< VaR_\alpha \\ &< -\widehat{\mu} - \sqrt{\widehat{\sigma^2}}\sqrt{(T-1)/q^{\chi^2_{T-1}}_{\frac{1-\beta}{2}}}\,q_\alpha \end{aligned} \tag{9}$$

where $q^{\chi^2_{T-1}}_{1-\beta/2}$ is the quantile of the chi-square distribution with $T-1$ degrees of freedom. (Note that one has slightly modified the original expression of Chappell and Dowd[4] by not setting μ to 0 as they did.) Their approach also induces that, when true parameters are known, the (exact) standard errors of the *VaR* estimate may be described (at a $\beta-$confidence level) by:

$$-\mu-\sigma\sqrt{q^{\chi^2_{T-1}}_{\frac{1-\beta}{2}}/(T-1)}q_\alpha < \widehat{VaR}_\alpha < -\mu-\sigma\sqrt{q^{\chi^2_{T-1}}_{\frac{1+\beta}{2}}/(T-1)}q_\alpha.$$

These last two expressions may be compared to those implied by the asymptotic results exposed in the third section.

For a large sample with size *T*, if μ is known, the variance of the (scaled) normal *VaR* estimate is asymptotically:

$$var\left(\sqrt{T}\ \widehat{VaR}_\alpha\right) = \frac{(VaR_\alpha+\mu)^2}{2} = \frac{1}{2}\sigma^2 q_\alpha^2.$$

Hence the associated (theoretical) standard error of the normal *VaR* estimate $S.E.\left(\widehat{VaR}\right)$ is $1/\sqrt{2T}|VaR+\mu|$ whereas related 'large sample' confidence intervals are

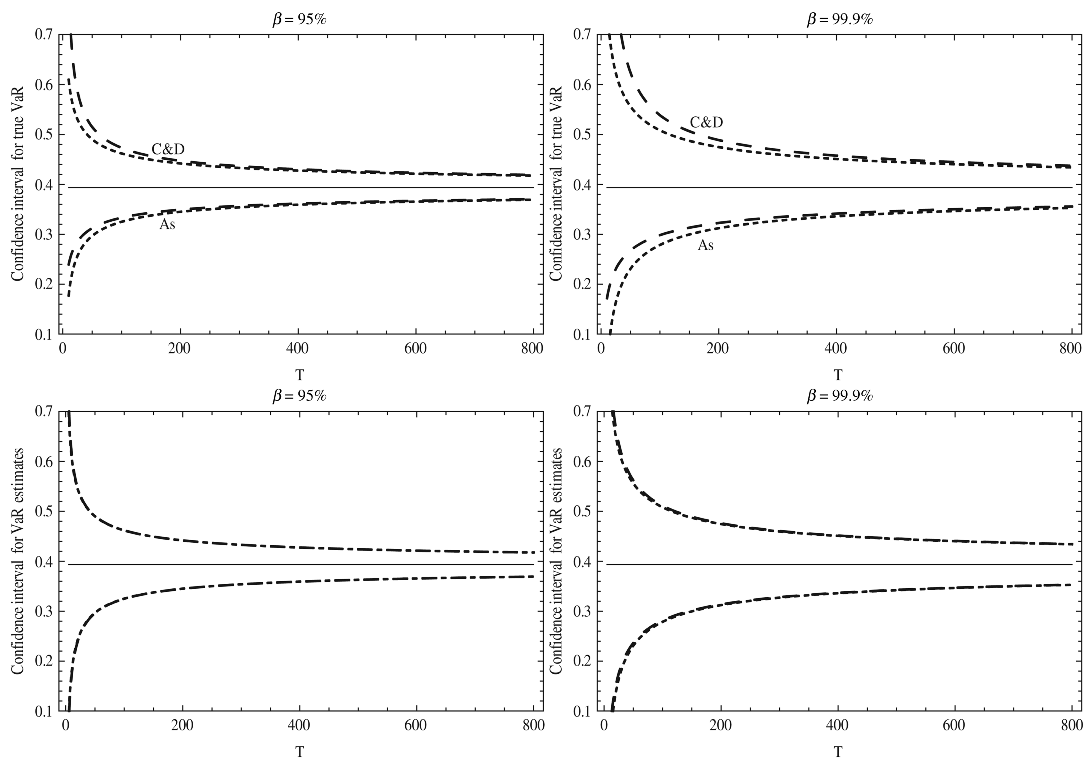

Figure 3: Assessing asymptotic confidence intervals
Notes: These graphs compare *asymptotic* confidence intervals for normal value-at-risk (*VaR*) with those deduced from the *exact* sampling distribution, given that the uncertainty about the mean has been neglected. The upper graphs plot confidence intervals for the true normal *VaR* given that $\widehat{\mu}$ and $\widehat{\sigma}$ are known. The lower graphs plot confidence intervals for the normal estimate $\widehat{VaR}$ given that μ and σ are known. The straight line stands for the centred value. The dashed line corresponds to the confidence interval deduced from the exact sampling setting while the dotted line stands for the asymptotic confidence interval. They are respectively labelled C&D and As.

given by:

$$-\widehat{\mu} - \sqrt{\widehat{\sigma^2}}\left(1 - \frac{1}{\sqrt{2T}} q_{\frac{1+\beta}{2}}\right) q_\alpha < VaR_\alpha < -\widehat{\mu} - \sqrt{\widehat{\sigma^2}} \times \left(1 - \frac{1}{\sqrt{2T}} q_{\frac{1-\beta}{2}}\right) q_\alpha \tag{10}$$

$$-\mu - \sigma\left(1 + \frac{1}{\sqrt{2T}} q_{\frac{1-\beta}{2}}\right) q_\alpha < \widehat{VaR}_\alpha < -\mu - \sigma\left(1 + \frac{1}{\sqrt{2T}} q_{\frac{1+\beta}{2}}\right) q_\alpha.$$

Note that, by construction, the asymptotic confidence intervals are symmetrical around the mean value. Figure 3 now investigates the large sample confidence intervals defined above.

The upper graphs of Figure 3 compare the asymptotic confidence intervals associated to the 'true' normal *VaR* (exposed in (10)) with those deduced from the Chappell and Dowd[4] approach (exposed in expression (9)). Here $\widehat{\mu}$ and $\widehat{\sigma^2}$ are known values. The lower graphs investigate expressions for normal *VaR* estimates and both μ and σ are supposed to be known. All these graphs show, for a given confidence level β, how the asymptotic confidence interval can approach the exact one. The lower graphs indicate that the asymptotic distribution derived in proposition 1 is an accurate approximation of the true one. The upper graphs recall that, using the exact *chi-square* distribution does not provide a symmetrical result confidence interval around the expected value. For moderately large samples, the asymptotic upper bound slightly underestimates the volatility uncertainty while the lower bound overestimates it. Depending on the confidence level, a set of one or two years of daily observations may be considered a large sample.

CONCLUSION

This paper suggests a new way to assess the estimation risk of *VaRs*. In short, one exploits the asymptotic theory and the so-called delta method to derive analytical results for obtaining confidence intervals of normal *VaR*. This paper extends the analytical approaches of Jorion[5,6] and Chappell and Dowd[4] by taking into account both the mean uncertainty and the volatility uncertainty. It avoids any recourse to simulations as in Dowd[7] or Cotter and Dowd.[8] Overall, the author believes this is a useful contribution to the general literature on *VaR*.

More qualitatively, the analytical results confirm the conclusions of Dowd[7] that the key factor behind the width of confidence intervals is sample size. To improve the estimation accuracy by a factor of two, it is found that four times as much data are needed (*ceteris paribus*). The estimation risk has been broken down into various components. It is shown, among other things, that the fraction of error owing to the mean uncertainty is of limited magnitude in a prudential context. In other words, the approximate approach defended by Jorion[5,6] and Chappell and Dowd[4] is shown to be especially relevant.

As a final word, it can be stressed that the paper also sheds light on a key issue of risk management. This is the best-practice or regulator-required number of data to be used. This paper shows that the standard error is around 5 per cent of the estimated *VaR* for 250 observations (and the usual confidence level) and that four times as much data are required to improve the estimation accuracy by a factor of two.

Acknowledgment

The author thanks Frances Maguire (JRMFI co-editor) and an anonymous referee for comments on the paper. Moraux's research is supported by the 'Fondation IGR-IAE' and the CREM (the Centre de Recherche en Economie et Management) CNRS unit (UMR CNRS 6211).

References

1 Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D. (1999) 'Coherent measures of risk', *Mathematical Finance*, Vol. 9, No. 3, pp. 203–228.
2 Christoffersen, P. and Gonçalves, S. (2005) 'Estimation risk in financial risk management', *Journal of Risk*, Vol. 7, No. 3, pp. 1–28.
3 Kendall, M. and Stuart, A. (1973) 'The Advanced Theory of Statistics. Volume 2: Inference and Relationship', 3rd edn, Griffin, London.
4 Chappell, D. and Dowd, K. (1999) 'Confidence intervals for VaR', *Financial Engineering News*, March, pp. 1–2.
5 Jorion, P. (1996) 'Risk^2: Measuring the risk in value at risk', *Financial Analysts Journal*, Vol. 52, No. 6, pp. 47–56.
6 Jorion, P. (2001) 'Value at Risk: The New Benchmark for Managing Financial Risk', 2nd edn, McGraw Hill, New York, NY.
7 Dowd, K. (2000) 'Assessing VaR accuracy', *Derivatives Quarterly*, Vol. 6, No. 3, pp. 61–63.
8 Cotter, J. and Dowd, K. (2007) 'Evaluating the precision of estimators of quantile-based risk measures', working paper, University College Dublin.

APPENDIX

This paper essentially accounts for the distribution of the normal *VaR* estimate. One could rather think to focus on the two first moments of the normal *VaR* estimate. Since $\widehat{VaR}_\alpha = -\widehat{\mu} - \widehat{\sigma} q_\alpha$, one may write:

$$E\left[\widehat{VaR}_\alpha\right] = -E[\widehat{\mu}] - q_\alpha E[\widehat{\sigma}] \qquad (11)$$

$$\mathrm{var}\left[\widehat{VaR}_\alpha\right] = \mathrm{var}[\widehat{\mu}] + q_\alpha^2 \mathrm{var}[\widehat{\sigma}] \qquad (12)$$

The expression (3) states that $\widehat{\mu}$ is an unbiased estimate of the true mean μ and that its variance is $\frac{\sigma^2}{T}$. Previous expressions can then be rewritten:

$$E\left[\widehat{VaR}_\alpha\right] = -\mu - q_\alpha E[\widehat{\sigma}],$$

$$\mathrm{var}\left[\widehat{VaR}_\alpha\right] = \frac{\sigma^2}{T} + q_\alpha^2 \mathrm{var}[\widehat{\sigma}].$$

The expression ($3'$) implies that the sample variance is unbiased ($E\left[\widehat{\sigma^2}\right] = \sigma^2$) and that its variance is $\mathrm{var}\left[\widehat{\sigma^2}\right] = 2\sigma^2/(T-1)$ (a χ^2_{T-1} random value has indeed a mean and a variance equal to $(T-1)$ and $2(T-1)$, respectively). And these results concern the unbiased sample variance only. This is a problem because computation of *VaR* requires results on $\widehat{\sigma} = \sqrt{\widehat{\sigma^2}}$. This associated sample estimate of the standard deviation is actually biased. Although the square root of the unbiased sample variance $\sqrt{\widehat{\sigma^2}}$ is a natural candidate to estimate σ, problems then arise because $E|\sqrt{\widehat{\sigma^2}}|$ is not necessarily equal to $\sqrt{E|\widehat{\sigma^2}|}$ (where $\sqrt{E|\widehat{\sigma^2}|}$ is of course equal to σ). Actually, the Jensen's

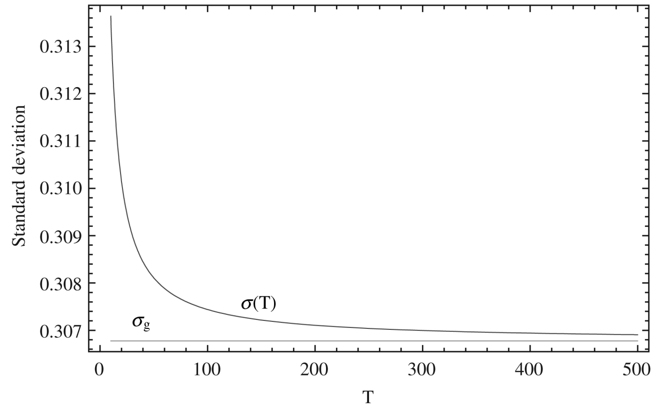

Figure A1: Standard deviation of $\sqrt{T}\widehat{VaR}$: 'analytical approximate' estimate vs asymptotic estimate
Note: This figure plots both the asymptotic estimate ($\sigma_g = \sqrt{V_g}$) and the analytical estimate $\sigma(T) = \sigma\sqrt{1 + q_\alpha^2(1 - c_4^2(T))\,T}$ of the standard deviation of $\sqrt{T}\widehat{VaR}$. One sets $\alpha = 95\%$ and $\sigma = 20\%$.

inequality even suggests that $E|\sqrt{\widehat{\sigma^2}}| \le \sqrt{E|\widehat{\sigma^2}|} \equiv \sigma$ because the square root is a concave function. $\sqrt{\widehat{\sigma^2}}$ is therefore a downwardly biased estimate of the (true) standard deviation σ. In other words, $\sqrt{\widehat{\sigma^2}}$ is (on average) too small! Statistical theory tells us what the bias is and it has been demonstrated that:

$$\frac{1}{c_4(T)} E[\sqrt{\widehat{\sigma^2}}] = \sigma$$

with

$$c_4(T) = \sqrt{\frac{2}{T-1}}\Gamma\left(\frac{T}{2}\right) / \Gamma\left(\frac{T-1}{2}\right)$$

where Γ is the Gamma function. In other words, $1/c_4(T)\sqrt{\widehat{\sigma^2}}$ is the unbiased sample estimate of the standard deviation. And, since $\mathrm{var}|\sqrt{\widehat{\sigma^2}}| = E|(\sqrt{\widehat{\sigma^2}})^2| - (E|\sqrt{\widehat{\sigma^2}}|)^2 = \sigma^2(1 - c_4^2(T))$, the standard error of the unbiased estimate of the sample standard deviation is $\sigma\sqrt{1 - c_4^2(T)/c_4^2(T)}$. Note that properties of the function $c_4(T)$ imply $c_4(T) \approx 1$, when T is very large, and $1 - c_4^2(T) \approx 1/2(T-1)$.

Given equations (11) and (12) and the above discussion, the analytical expressions for the mean and the variance of the normal *VaR* estimate become:

$$E\left[\widehat{VaR}_\alpha\right] = -\mu - q_\alpha\sqrt{\frac{2}{T-1}}\frac{\Gamma\left(\frac{T}{2}\right)}{\Gamma\left(\frac{T-1}{2}\right)}\sigma \underset{\text{if } T \text{ is large}}{\approx} -\mu - q_\alpha\sigma = VaR$$

$$\mathrm{var}\left[\widehat{VaR}_\alpha\right] = \frac{\sigma^2}{T} + q_\alpha^2\left(1 - c_4^2(T)\right)\sigma^2 \underset{\text{if } T \text{ is large}}{\approx} \frac{\sigma^2}{T}\left(1 + \frac{1}{2}q_\alpha^2\frac{T}{T-1}\right)$$

The first equation confirms that (for large T) the $\widehat{VaR}_\alpha$ is an unbiased estimate of the true *VaR*. The second one implies (again for large T) that the variance is:

$$\mathrm{var}\left[\sqrt{T}\ \widehat{VaR}_{\alpha}\right] \underset{\text{if } T \text{ is large}}{\approx} \sigma^2 \times \left(1 + \frac{1}{2} q_{\alpha}^2 \frac{T}{T-1}\right).$$

This analytical approximation is directly comparable to the asymptotic variance V_g exposed in proposition 1. But this approach provides no confidence intervals for the *VaR* estimate. For completeness, Figure A1 displays the convergence of $\sqrt{\mathrm{var}\left[\sqrt{T}\ \widehat{VaR}_{\alpha}\right]}$ to $\sqrt{V_g}$.

