The Cost of Conservatism:

Extreme Returns, Value-at-Risk, and the Basle 'Multiplication Factor'[*]

Jon Danielsson
London School of Economics
and IoES at University of Iceland

Philipp Hartmann
London School of Economics

Casper G. de Vries
Erasmus University Rotterdam and the Tinbergen Institute

This paper was published in the January 1998 issue of Risk

Abstract

We argue that most current methodologies for value-at-risk (VaR) under-estimate the VaR, and are therefore ill-suited for market risk capital. Better VaR methods are available, such as the tail–£tting method proposed here. However, £nancial institutions may be reluctant to use those methods since current VaR regulation may, perversely, provide incentives for banks to under-estimate VaR as much as possible.

[*]Email: *j.danielsson@lse.ac.uk*, *p.hartmann@lse.ac.uk*, and *cdevries@few.eur.nl*. The authors research papers can be downloaded from the web site *www.RiskResearch.org* or *http://fmg.lse.ac.uk/~jond.*

1 Introduction

In January 1998 the amendment for market risks of the Basle Capital Accord will become effective in the G-10 countries. The main change, compared to previous capital adequacy regulations, is the option for banks to use their own internal market risk management models, i.e. value-at-risk (VaR) models from which their regulatory minimum capital against trading book losses is determined. In this article we point to several important facts, which we feel have been neglected in the discussion about VaR models in general and the Basle internal models approach in particular. We argue that the current set of Basle requirements still provides disincentives for the development of more reliable VaR models, and show that considerable improvement of current VaR models is possible by means of techniques that explicitly focus on the properties of extreme return ¤uctuations. We then brie¤y discuss how a change in the determination of the Basle 'multiplication factor' may encourage the industry to adopt improved VaR models, such as those proposed here.

2 Traditional VaR models and extreme returns

VaR models usually use historical data to evaluate maximum (worst case) trading losses for a given portfolio over a certain holding period at a given con£dence interval. For example, a VaR model may tell you that a banks' daily trading loss, of 1 million dollars or more, will occur with 3% probability. A £rst important observation is that value-at-risk applies to the extreme lower tail of the return frequency distribution, i.e. large losses, far away from the mean. This fact is recognized in the Basle market risk amendment, which speci£es the use of a one-sided con£dence interval of 99 percent, i.e. the chance of experiencing a larger loss than the value at risk should be 1 in 100 or less.

This number obviously re¤ects regulators' natural tendency for conservativeness in their prudential supervision of banks. The same tendency also comes out in the Basle regulators' choice of holding period, a second important model parameter. While the industry virtually unanimously uses daily VaRs for internal risk control, for the purpose of determining their minimum regulatory capital against market risk, banks will be obliged to assume that they cannot liquidate their trading portfolios quicker than within 10 business days. In order to facilitate the transition from their internal daily VaR models to the regulatory 10-day models the application of a simple 'square-root-of-time' rule is permitted. We shall come back to this rule below, making the argument that - surprisingly - this may lead to an over-estimation of the bi-weekly VaR.

The most dif£cult part in VaR estimations is the derivation of the portfo-

lio return frequency distribution. Two approaches have become widely popular: Variance-covariance analysis and historical simulation. Variance-covariance analysis relies on the assumption that £nancial market returns follow a multivariate normal distribution. It is easy to implement, because the VaR can be computed from a simple linear formula with variances and covariances of returns as the only inputs. Its major drawback is that £nancial market returns are not normally distributed, having fatter tails than the normal. This means that losses are much more frequent than predicted by variance-covariance analysis. It is particularly weak where a VaR model for regulatory purposes and risk control should be strong, i.e. in the prediction of large losses.

Another feature of many VaR models is the exponential weighing of past returns, i.e. returns closer to the present are given more weight than those several months or even a year ago. This technique is justi£ed by the presence of conditional heteroscedasticity (CH) in daily £nancial market returns, meaning that a volatile day is usually followed by volatile days. However, two important observations are relevant for the Basle minimum capital requirement discussion. First, while daily returns exhibit strong CH effects, they can hardly be detected in biweekly returns such as the regulatory 10-day holding period. Second, CH effects largely originate from medium and small range volatility periods. Extreme events, such as losses at or beyond a 99% con£dence interval, scatter rather independently over time Daníelsson and de Vries (2000). Jackson, Maude, and Perraudin (1997).

Historical simulation does not suffer from the tail-bias problem, because it does not rely on normality. By applying the full empirical market return distribution to all the items in the current trading portfolio, the outcome exactly re¤ects the historical frequency of large losses over the speci£c data window. Another advantage of this approach compared to variance-covariance analysis is that it can incorporate non-linear positions, such as derivative positions, in a natural way; see Kupiec and O'Brien (1995a) on this "full-valuation" property. The problem with historical simulation is that it is very sensitive to the particular data window, which the Basle Committee has chosen to be at least one year of past returns. In other words, whether October 1987 is included or not makes a huge difference for the value at risk predicted. Stated differently, the empirical return distribution is very 'dense' and smooth around the mean, so that no parametric model based on a standard distribution, such as the normal, can beat the accuracy of the empirical distribution there. Due to the few occurrences of extremely large price movements, however, it becomes 'discrete' in the tails. Hence, VaR predictions based on historical simulation exhibit high variances. Moreover, at its lower end, the empirical return distribution sharply drops to zero and remains there, i.e. more severe losses in the future than the largest one during the past year is given a probability of zero, which might be considered as imprudent.

Our interim conclusion is that a good value-at-risk model to satisfy regulatory

minimum capital standards should correctly represent the likelihood of extreme events by providing smooth tail estimates of the portfolio return distribution which extend beyond the historical sample (up to in£nity). Exponential or other weighing schemes need not be adopted. In what follows we shall sketch a new VaR model which satis£es these requirements by combining historical simulation for the interior of the portfolio return distribution with a parametric estimator for the tails.

3 A new approach based on extreme-value theory

Our predictions for tail events will be based on the well documented fact that asset return data have fatter tails than the normal. All heavy-tailed distributions eventually display the same tail behavior when we consider large losses or gains. Heavy-tailed distributions all have tail shapes which to a £rst order approximation are identical to the tail shape of the Pareto distribution. Thus if we know that the distribution is heavy tailed, then the largest losses occur with the following approximate probability:

$$\Pr\{R < -r\} = F(-r) \approx ar^{-\alpha}, \ \text{ as } r \to \infty \tag{1}$$

The tail probabilities depend on two parameters, a scaling constant a, and the so called tail index α. (1) shows that the smaller the tail index is, the more likely are extreme events and the fatter are the tails. Statistical extreme value theory provides the tools for estimating these tail parameters. Daníelsson and de Vries (2000) and Daníelsson and de Vries (1997) for a more detailed description of the theory and applications.) We have estimated these parameters for a number of return data series.

In two interesting contributions in this journal, Boudoukh, Richardson, and Whitelaw (1995) and Bahar, Gold, and Pilizu (1997) argue that VaR should focus on the worst case scenario by means of studying the distribution of the minimum return out of a number of n returns. This distribution is known to be $-F[1 - F(-r)]^n$, and can be computed if the underlying distribution $F(-r)$ is given. Interestingly, what these articles did not recognize is, that if n becomes large, then the distribution of the minimum converges to a limit distribution which is known a priori (the so-called extreme value distribution). If the data are fat tailed then the leading term in a Taylor expansion of this extreme value distribution is given by (1). Thus our approach genuinely takes care of the worst case scenario, but does not require prior knowledge of the speci£c distribution.

In order to illustrate the differences between the tail estimation technique and other methods for VaR inference, such as variance-covariance analysis and pure historical simulation, a particularly volatile asset class is used, i.e. daily returns

on spot oil prices from 1986 to 1997. By applying the method of Daníelsson and de Vries (1997) we can predict from the derived tail probabilities that the maximum expected one-day drop in oil prices during a period of 15 years is 28%, and that a drop of 44% is expected once every 70 years. In a single year we expect one day where prices increase by 10%, on average.

Figure 1: Lower Tail of Daily Oil Prices 1992-1997
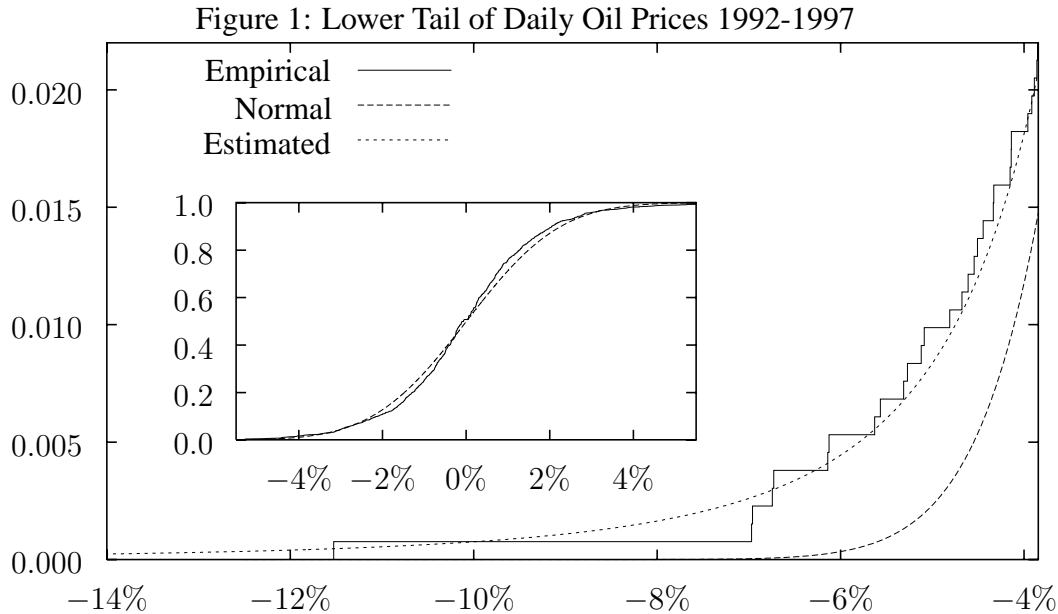
The Figure illustrates a subsample of the results from 1992 to 1997. It shows the results from three different techniques for VaR estimation; historical simulation, variance-covariance methods, and extreme-value analysis. The step function gives the empirically realized losses on the horizontal axis with the frequency of losses on the vertical axis and reﬂects the results one would obtain from using historical simulation. (The smaller window gives the entire empirical cumulative distribution function.) The dotted line plots the estimated tail probabilities. The Figure clearly illustrates how the tail estimator smoothes the distribution of extreme returns. Moreover, the curve extends 'beyond the sample' (to the left), allowing evaluation of the frequency and magnitude of losses for a much larger time interval than we have data for. It is instructive to compare this with the losses predicted by variance-covariance analysis, which is indicated by the dashed line. This curve is located way below the other two, in the southeast corner. Clearly, the normality assumption leads to a substantial under-estimation of VaR. For example, the normal distribution predicts a 13 times lower probability for a 6% drop of the oil price than the estimated tail distribution.

Daníelsson and de Vries (2000) compare the performance of various VaR prediction methods for simulated portfolios of US stocks over a 4-year period. A

5

Table 1: Performance of Different VaR Models (1000 day horizon)

Con£dence level	95%	99%	99.95%
Expected number of VaR violations	50	10	0.5
Average number of realized VaR violations (percentage error in parenthesis)			
Variance-Covariance Approach	52.45	16.28	3.55
	(4%)	(63%)	(610%)
Historical Simulation	43.24	7.66	0
	(-14%)	(-23%)	(-100%)
Tail Estimator	43.14	8.19	0.59
	(-14%)	(-18%)	(18%)

Note: Models tested with 500 random portfolios comprising 7 US stocks for the period 1993 through 1996.

subset of the results is shown in the Table, which compares the expected number of VaR violations for three different con£dence intervals with those actually realized in the simulations. For example, for the Basle 99% con£dence interval, 10 violations are to be expected; the variance-covariance approach gave on average 16 violations, while for the tail estimation approach the actual portfolio loss exceeded the predicted value at risk in only 8 cases. For the 95% level the variance-covariance method performs best with 4% error, but as the con£dence level increases the performance of this method steadily worsens, e.g. at the 99% con£dence level the error is over 60%. Historical simulation has mixed performance; it provides reasonable predictions at the lower con£dence levels, but it worsens with higher levels and becomes uninformative at the out-of-sample 99.95% con£dence level. Note that the tail estimator still provides a remarkably good estimate at that level.

4 Surprises from time aggregation

As has been pointed out before Kupiec and O'Brien (1995a), the simple 'square-root-of-time' formula to aggregate daily VaR estimates to bi-weekly estimates can be rather imprecise when returns of the underlying market risk factors are non-normal. This has led to concerns among regulators that the simple time-aggregation rule can easily lead to an under-estimation of potential losses and therefore to too little capital against market risk, suggesting a more stringent regulatory approach.

We rather come to the opposite conclusion and argue that the square-root formula may lead to an over-estimation of value at risk, when returns are not nor-

mally distributed and exhibit fat tails. Assume that observed £nancial market returns r have £nite variance. This implies that the tail index α is larger than 2 in (1). Let t denote the length of the holding period. Increasing the holding period t increases the VaR under the normal model by a factor of $\sqrt[2]{2}$('square-root-of-time rule'). If the return distribution is fat tailed, then this factor equals $\sqrt[\alpha]{2}$, and since $\alpha > 2$, $\sqrt[2]{2} > \sqrt[\alpha]{2}$. Hence the 'square-root-of-time' rule eventually results in a higher VaR than the value implied by a heavy tailed distribution. This result follows from equation (1) and is a direct implication of the linear tail additivity of fat tailed distributed random variables, and the self-additivity of normally distributed variables. (See e.g. Dacorogna, Muller, Pictet, and de Vries (1999) for further details.)

In sum, by prescribing the 'square-root-of-time' rule for time-aggregation regulators have - consciously or unconsciously - introduced another element of conservativeness in the internal models approach to market risk capital requirements, which has passed widely unnoticed.

5 Incentives and the Basle multiplication factor

The regulatory requirement that banks' 10-day, 99 % VaR estimate has to be multiplied by a factor of at least 3 to determine the minimum regulatory capital against market risk has received a cool reception by the industry. National bankers associations argued that such a high factor would discourage the application of quantitative models and obstruct progress in risk management techniques Elder£eld (1995). For example, even when applying the tail estimation approach proposed above, which is clearly much more precise than more widely known standard approaches, a bank could not be granted a lower factor than 3. Nevertheless, the Basle Committee has con£rmed in December last year that it will retain the size of the factor. In addition, this factor can be increased through a variable add-on between 0 and 1 depending on the performance of a bank's 1-day model in backtesting procedures. The Basle Committee points out that the variable component provides for built-in incentives to develop and use better models. However, we would argue that the £xed component of 3 is already so high that it completely dominates any potential advantages from achieving a zero add-on through a good model. In fact, we would propose that the £xed component should be lowered and the range of the variable add-on potentially extended in order to leave suf£cient incentives for banks to use the best models.

Stahl (1997) has recently advanced an interesting theoretical justi£cation for the £xed factor of 3, which has formerly been interpreted as a somewhat arbitrary political compromise. His two arguments, one related to the tail misspeci£cation in variance-covariance approaches, the other related to potential time-variation of

portfolio return distributions over the relevant data window (in particular uniden-ti£ed increases in the variance of returns), are both based on a very general statis-tical result known as Chebyshev Inequality. If R is a random variable with mean μ_R and £nite variance σ_R^2, then

$$\Pr\left[\mu_R - k\sigma_R < R < \mu_R + k\sigma_R\right] \geq 1 - \frac{1}{k^2} \tag{2}$$

Equation (2) implies, for example, that – whatever the true distribution of a random variable R – the boundaries of a 99% con£dence interval ($\Pr\left[\cdot\right] = 0.99$) are never wider than 10 standard deviations left and right from the mean. $\left(0.99 \geq 1 - 1/k^2 \Leftrightarrow |k| \leq 10\right)$

Note that Chebyshev's inequality (2) is true for any type of distribution which has a £nite variance. Applied to VaR, it says in a way: Assume we do not know anything about the structure of £nancial market returns, what is the ex-treme boundary that could cover any speci£cation error? Of course, at this level of generality the answer must be a very conservative multiplication factor (some-thing between 3 and 4 according to Stahl's calculations) in order to cover even the weirdest distributions. An example showing how far off this theoretical bound can be, is given in Haan, Jansen, Koedijk, and de Vries (1994).

We do know much more than 'nothing' about £nancial market returns. As discussed above, we know that these returns have fat tails and that a single limit law determines the shape of these tails. Including this information, as done in the tail estimation approach described above, is both ef£cient from the risk manager's perspective and prudent from the regulator's point of view.

This conclusion is not subjective and hardly dependent on any particular em-pirical speci£cations. The £xed component of 3 in the Basle multiplication fac-tor is unnecessarily conservative. In order to give banks the opportunity to reap the bene£ts of better VaR models, in terms of lower minimum regulatory capital against market risks, the £xed component of the Basle multiplication factor should be reduced. In order to suf£ciently penalize bad models the range of the variable add-on could be increased. This would avoid any disincentives through pruden-tial capital requirements to future progress in banks' risk management techniques, while preserving the fundamentals of the Basle internal models approach Good-hart, Hartmann, Llewellyn, Rojas-Suarez, and Weisbrod (1997) (chapter 5).

To be sure, it is extremely dif£cult for external regulators to evaluate to which extent they can have con£dence in banks' internal risk management techniques, and some degree of conservativeness is a necessary characteristic of every £nan-cial regulator confronted with the sometimes wild gyrations of £nancial market behavior. But it seems that in the "Amendment to the Basle Capital Accord to Incorporate Market Risks" this conservativeness may have gone too far.

6 Conclusions

In this article we have argued that most traditional VaR techniques fail to properly model the tails of portfolio return distributions - the very essence of the value-at-risk concept. A semi-parametric approach, based on results from extreme value theory, turns out to produce much more reliable VaR estimates. Furthermore we argue that, in the light of theoretical and empirical considerations, the £xed Basle multiplication factor of 3 appears to be very conservative and should be reduced, while the range of the variable add-on should be increased.

There are already more far-reaching reform proposals on the table. Not the least the pre-commitment approach for market risk capital requirements Kupiec and O'Brien (1995b) and Kupiec and O'Brien (1997). Although rarely pointed out, pre-commitment actually implies an endogenous, incentive compatible multiplication factor. These new ideas have their merits, but before they are ready for political compromise, the present regulatory approaches should be improved to ensure £nancial stability at reasonable costs for both consumers and market participants.

References

BAHAR, R., M. GOLD, AND C. PILIZU (1997): "Making the Best of the worst," *Risk*, August, 100–103.

BASEL COMMITTEE ON BANKING SUPERVISION (1996): *Amendment to the Capital Accord to Incorporate Market Risks*.

BASLE COMMITTEE ON BANKING SUPERVISION (1996): *Supervisory Framework for the Use of "Backtesting" in Conjunction with the Internal Models Approach to Market Risk Capital Requirements*.

BOUDOUKH, J., M. RICHARDSON, AND R. WHITELAW (1995): "Expect the worst," *Risk*, September, 100–101.

DACOROGNA, M. M., U. A. MULLER, O. V. PICTET, AND C. G. DE VRIES (1999): "Extremal foreign exchange rate returns in extremely large data sets," *www.few.eur.nl/few/people/cdevries/workingpapers/workingpapers.htm*.

DANÍELSSON, J., AND C. G. DE VRIES (1997): "Tail index and quantile estimation with very high frequency data," *Journal of Empirical Finance*, 4, 241–257.

DANÍELSSON, J., AND C. G. DE VRIES (2000): "Value at Risk and Extreme Returns," London School of Economics, Financial Markets Group Discussion Paper, no. 273, *www.RiskResearch.org*.

ELDERFIELD, M. (1995): "Capital Incentives," *Risk*, September, 20f.

GOODHART, C. A. E., P. HARTMANN, D. T. LLEWELLYN, L. ROJAS-SUAREZ, AND S. R. WEISBROD (1997):. Monograph prepared for the Bank of England Central Bank Governors' Symposium, forthcoming with Routledge.

HAAN, L., D. W. JANSEN, K. KOEDIJK, AND C. G. DE VRIES (1994): "Safety £rst portfolio selection, extreme value theory and long run asset risks," pp. 471–487. Kluwer Press.

HARTMANN, P. (1997): "Capital Adequacy and Foreign Exchange Risk Regulation: Theoretical Considerations and Recent Developments in Industrial Countries," *Kredit und Kapital*, 30, 186–218.

JACKSON, P., D. J. MAUDE, AND W. PERRAUDIN (1997): "Bank capital and Value-at-Risk," *Journal of Derivatives*, Spring 1997, 73–111.

KUPIEC, P. M., AND J. M. O'BRIEN (1995a): "Internal Affairs," *Risk*, May, 43–48.

———— (1995b): "Model Alternative," *Risk*, June, 37–41.

———— (1997): "The Pre-Commitment Approach: Using Incentives to Set Market Risk Capital Requirements," FEDS Paper, no.97-14, Federal Reserve Board.

STAHL, G. (1997): "Three cheers," *Risk*, 10, 67–69.