Long-Term Value at Risk

By

Kevin Dowd, David Blake, and Andrew Cairns [*]

Revised, September 2003

[*] Dowd is at the Centre for Risk and Insurance Studies, Nottingham University Business School; Blake is at the Pensions Institute, Birkbeck College, University of London; and Cairns is at the Department of Actuarial Mathematics and Statistics, Heriot-Watt University. The authors thank Peter Simmons of UBS Global Asset Management and various anonymous referees whose comments have improved the paper; however, the usual caveat applies.

The corresponding author is: Kevin Dowd, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK; email: Kevin.Dowd@Nottingham.ac.uk.

Long-Term Value at Risk

Abstract

This paper investigates the estimation of long-term VaR. It also suggests a simple approach to the estimation of long-term VaR that avoids problems associated with the square-root rule for extrapolating VaR, as well as those associated with attempts to extrapolate day-to-day volatility forecasts over longer horizons.

One of the most significant recent developments in the risk measurement and management area has been the emergence of Value-at-Risk (VaR). The VaR of a portfolio is the maximum loss that the portfolio will suffer over a defined time horizon, at a specified level of probability known as the VaR confidence level. The VaR has proven to be a very useful measure of market risk, and is widely used in the securities and derivatives sectors: a good example is the RiskMetrics system developed by J. P. Morgan. VaR measures based on systems such as RiskMetrics' sister, CreditMetrics, have also shown their worth as measures of credit risk, and for dealing with credit-related derivatives. In addition, VaR can be used to measure cashflow risks and even operational risks.[1] However, these areas are mainly concerned with risks over a relatively short time horizon, and VaR has had a more limited impact so far on the insurance[2] and pensions literatures[3] that are mainly concerned with longer-term risks.

Yet the VaR literature also has relatively little to say on longer-term risk

---

[1] For more on VaR and its applications, see, e.g., Dowd (2002) or Barry Schachter's website on VaR, www.gloriamundi.org.

[2] A notable exception is an article by Panning (1999), which applies VaR to property/casualty insurers. The Panning article deals with four main issues: estimation risk, the impact of a changing portfolio, franchise risk, and the application of VaR to long-term risk management. By contrast, our paper focuses on only one issue (i.e., the estimation of long-term VaR) and covers this issue more comprehensively than Panning does (e.g., it examines the effects of the time horizon more closely, and has more to say on subsidiary issues such as volatility estimation).

[3] However, the issues involved in VaR are clearly related to the issues that arise in the probability-of-ruin literature, and there have been some attempts to apply VaR techniques to pension funds (e.g., the

measurement. Perhaps the best-known advice it offers is the square-root rule, and even that is usually applied to short time horizons. If VaR(h) is the VaR over a horizon of h days, and VaR(1) is the VaR over one day, this rule tells us that we can obtain the former from the latter by multiplying it by the square root of h:

$$VaR(h) = VaR(1)\sqrt{h} \tag{1}$$

Such scaling is widely used, and is enshrined prominently in the Market Risk Amendment to the Basle Accord.[4] Unfortunately, this rule is unreliable, and can lead to considerable overestimates of VaR (see, e.g., Blake *et alia* (2000)). There also seems to be a general feeling among practitioners that the estimation of longer-term VaR is more difficult than the estimation of short-term VaR. This perception owes much to problems of longer-term volatility forecasting, the argument being that VaRs depend on volatility, and volatility is (much) more difficult to forecast over longer horizons (e.g., Christoffersen *et al*. (1998, p. 109)).

This paper offers a different approach to this problem. Our approach goes back to first principles and suggests that the estimation of long-term VaR is actually quite straightforward. The idea is to apply a standard quantile formula over the long-term horizon, and then estimate VaR using estimates of the horizon-average values of

PensionMetrics approach of Blake *et al.* (2001), or Gupta *et al*. (2000))

[4] Specifically, the Market Risk Amendment suggests that banks should estimate VaR for a 10-day horizon, and banks are allowed to obtain these estimates by scaling up shorter-horizon VaRs using the square root rule (Basle Committee (1996, Section B.4, paragraph c, p. 44)).

the parameters on which the VaR depends. This approach does not require us to forecast day-to-day volatilities over long horizons, and so avoids the (real) difficulties of standard volatility-forecasting approaches. We also suggest that the estimation of long-term VaR should *not* involve the square-root rule, which can be misleading, even for relatively short horizons, and is especially misleading for longer ones.

The outline of this paper is as follows. Section 1 provides the basic analytical framework. Section 2 then looks at how VaR varies with the holding period, and section 3 carries out some sensitivity analysis and, in particular, looks at the sensitivity of VaR estimates to changes in the mean and volatility of returns. Section 4 discusses the derivation of the return parameters for long-term VaR, and suggests that extrapolating traditional day-to-day forecast techniques is effectively useless in this context. Instead, the best approach is simply to take a view about the values of the mean long-term parameters involved. Some conclusions are offered in section 5.

1. Basic Analysis

Suppose we have a portfolio that generates a random daily real log-return with mean $\mu > 0$ and standard deviation (or volatility) σ. Positive return observations correspond to profits, and negative ones to losses, and we assume for convenience that any interim profits/losses are ploughed back into the portfolio and that the composition of our portfolio does not change over our investment horizon.[5] The VaR confidence level is cl and we consider VaR over a horizon of h days.

To illustrate the method, assume that daily log-returns are normally distributed. This lognormal assumption is very convenient for VaR analysis, but we also get similar results if we make the alternative assumption that log-returns are Student-*t* distributed.[6] The VaR associated with normally distributed log-returns is:

$$VaR(h) = P - P_{cl} = P - \exp[\mu h + \alpha_{cl}\sigma\sqrt{h} + \ln P] \quad (2)$$

where is the current value of our portfolio, is the (1-) percentile (or critical h

h

so α_{cl} = -1.645 if we have a 95% confidence level; see, e.g., Dowd (2002, p. 43)).

2. VaR and Time Horizon

We now consider how the VaR alters with the time horizon. A typical example is shown in Figure 1, based on annualized parameter values of $\mu = 0.075$ and $\sigma = 0.25$ and an initial portfolio value of $1. This illustrates how VaR changes with both time horizon and confidence level. For any given confidence level, as the time horizon increases, the VaR rises initially but then peaks and turns down; after that it keeps falling, becomes negative at some point, and thereafter remains negative and moves further and further away from zero.[7] The behaviour of the VaR also depends on the confidence level: for relatively low confidence levels, the VaR peaks quickly and then rapidly falls; but for relatively high confidence levels, the VaR peaks slowly and stays at or near its maximum value – which is bounded above by, and sometimes close to, the value of the investment itself – for a long time. Note, too, that whilst the VaR has this natural upper bound, it has no corresponding lower bound, and will fall indefinitely as the horizon continues to rise.

[7] A negative VaR simply means that the likely worst outcome at the specified level of confidence is a profit, rather than a loss.

Figure 1: VaR and the Time Horizon

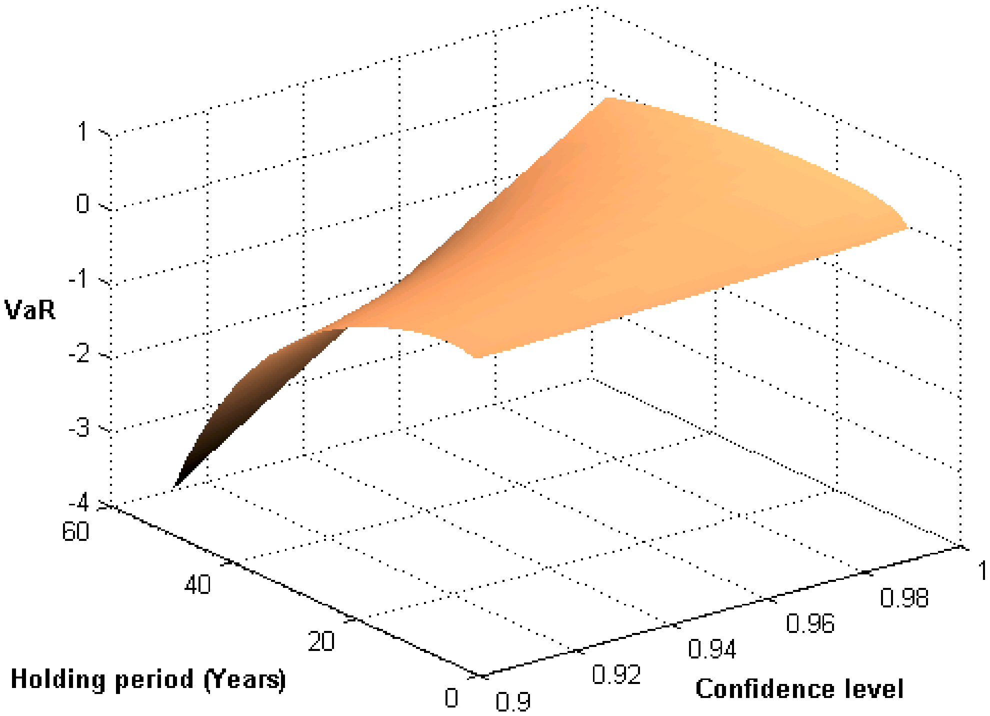

Note: Based on assumed parameter values of $\mu = 0.075$ and $\sigma = 0.25$, and an assumed initial investment of $1.

The VaR surface always retains this same shape provided that μ and σ are both positive. However, if $\mu = 0$, the VaR surface takes the rather different shape shown in Figure 2: the VaR approaches its ceiling asymptotically, and stays in that region indefinitely; and it approaches this maximum more quickly for the higher confidence levels. The story is therefore obvious: the VaR will initially rise, and will rise to its maximum possible value; however, when $\mu > 0$, the compounding of the mean return over the time horizon will eventually bring it down, and it will continue to fall thereafter.

Figure 2: VaR and Time Horizon with a Zero Mean Return

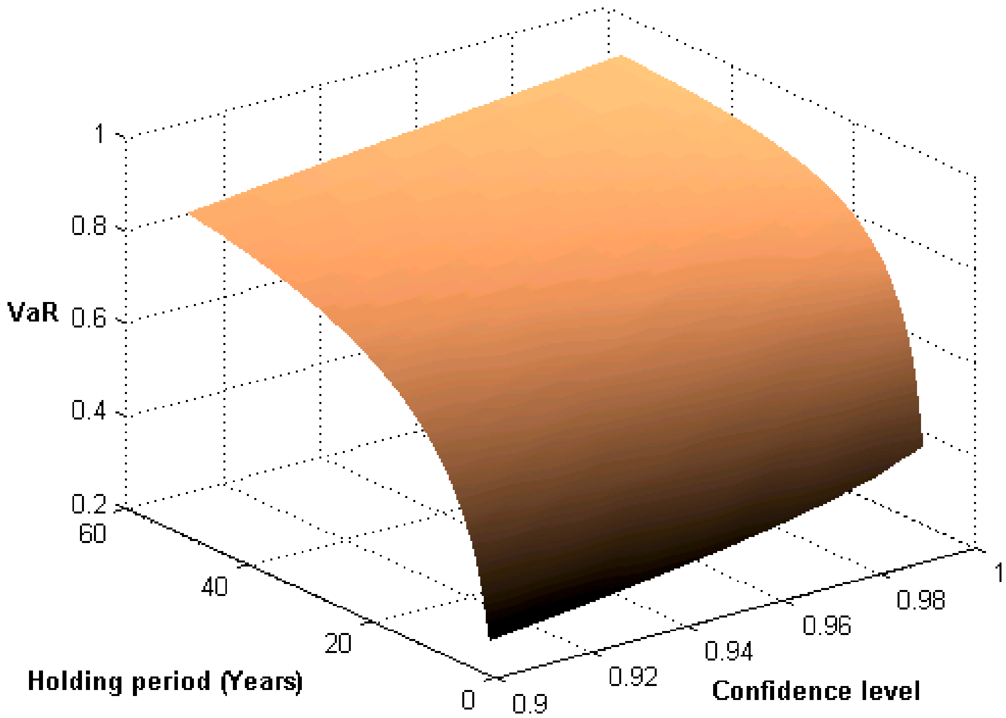

Note: Based on assumed parameter values of $\mu = 0$ and $\sigma = 0.25$, and an assumed initial investment example, the VaR at the 95% confidence level rises to 0.830 when the horizon reaches 20 years, and is still increasing with the horizon.

- If μ is relatively high and σ relatively low, the VaR rises and then falls relatively quickly. The comparable VaR – at the 95% confidence level and 20-year horizon period – in this case has already peaked and fallen to -1.451, and continues to fall with the time horizon.

Table 1: VaR and Time Horizon

Horizon (years)	**1**	**2.5**	**5**	**10**	**20**	**40**
			Low μ, high σ			
VaR at 95% *cl*	0.415	0.555	0.663	0.758	0.830	0.870
VaR at 99% *cl*	0.539	0.695	0.802	0.886	0.942	0.971
			High μ, low σ			
VaR at 95% *cl*	0.137	0.131	0.050	-0.246	-1.451	-10.468
VaR at 99% *cl*	0.220	0.261	0.244	0.098	-0.552	-5.010

Note: Figures are VaRs based equation (1), an initial investment of $1, and assumed parameter values of $\mu_{low} = 0.04$, $\mu_{high} = 0.10$, $\sigma_{low} = 0.15$ and $\sigma_{high} = 0.35$.

It is also clear that the square-root VaR will generally be very inaccurate over longer periods. Equation (1) indicates that the square-root VaR will rise indefinitely, proportionately to the square root of the time horizon, if we make the reasonable assumption that the initial, one-day VaR, is positive. At some point, it will therefore break through the VaR's (usual) natural upper barrier – the value of the investment –

and grossly over-estimate the VaR. By contrast, the true VaR will rise toward the barrier and then fall again, and ought never to exceed the value of the investment given limited liability. The magnitude of the error associated with the square-root rule thus rises with the time horizon. In addition, this error rises with μ, because the square-root formula makes no proper allowance for the impact of the compounding of μ in the VaR.[8]

3. Sensitivity Analysis

Sensitivity of VaR estimates to mean return

The next stage in our analysis is to examine the sensitivity of our VaR estimates to changes in various assumptions, and we begin by looking at their sensitivity to mean returns. To do so, we increase the assumed daily mean return by 1% of its value, and derive the associated percentage change in VaRs. Our results indicate that the sensitivity of VaR to estimated mean return is generally low over short time horizons. However, the sensitivity of our VaRs to the estimated mean also tends to rise in

[8] It is clear from the VaR equation that this result depends in part on the assumption that $\mu > 0$, and that for any given *h*, the degree of over-estimation increases directly with μ. We would also argue that the assumption that $\mu > 0$ is not unreasonable if we are considering investments, although it might be problematic in some insurance contexts (e.g., dealing with loss reserves). However, we would emphasize that the basic VaR approach is not contingent on any particular assumptions about the mean return, and we can easily estimate VaRs assuming zero or negative mean returns if we ever wanted to.

absolute terms, and eventually changes sign. Thus, broadly speaking, the VaR estimates become more sensitive to assumed mean returns, the longer the time horizon on which the VaRs are based.

Table 2: Sensitivity of VaR to Mean Return

	Percentage Change in VaR					
Horizon (Years) =	**1**	**2.5**	**5**	**10**	**20**	**40**
% Change in VaR at 95% *cl*	-0.2%	-0.3%	-0.5%	-1.0%	-3.8%	9.0%
% Change in VaR at 99% *cl*	-0.1%	-0.2%	-0.3%	-0.4%	-0.8%	-3.0%

Note: Based on assumed parameter values of $\mu = 0.075$ and $\sigma = 0.25$, and a +1% change in μ.

Sensitivity of VaR estimates to return volatility

We now look at the sensitivity of VaR to the volatility of returns. Table 3 reports some illustrative results showing percentage changes in VaR conditional on a 1% increase in volatility. Generally speaking, we tend to find that the sensitivity of the VaR to volatility increases with the holding period and, at least for low confidence levels, eventually changes sign as well. These results show that VaR is sensitive to volatility assumptions, and that the effect of a change in volatility on VaR depends importantly on the length of the time horizon[9].

[9] The length of the time horizon also influences the way in which volatility is estimated. For example,

Table 3: Sensitivity of VaR to Return Volatility

	Percentage Change in VaR					
Horizon (Years) =	**1**	**2.5**	**5**	**10**	**20**	**40**
% Change in VaR at 95% *cl*	1.0%	1.1%	1.3%	1.8%	4.5%	-7.8%
% Change in VaR at 99% *cl*	0.9%	0.9%	0.9%	0.9%	1.3%	3.7%

Note: Based on assumed parameter values of $\mu = 0.075$ and $\sigma = 0.25$, and a +1% change in σ.

4. Deriving the Return Parameters for Long-Term VaR

We have assumed so far that we already have estimates of the mean and volatility of returns that apply over our time horizon. But how do we derive these?

One approach is to forecast them using conventional forecasting methods. We could break up our horizon into a series of successive sub-periods (e.g., days) and forecast our mean return or volatility for each day in our horizon period. We could then use these forecasts to construct an estimate of the mean return or volatility for our whole horizon period or feed them into a more complex multi-period VaR analysis (e.g., such as a Monte Carlo simulation).

it is reasonable to estimate monthly volatility using daily data. But it would not be sensible to estimate annual volatility using daily data: a more reliable estimate would be based on monthly data.

Unfortunately, this approach runs into various difficulties. One problem is that any long-run forecasting rule will eventually give implausible forecasts if the variable being forecasted has a trend and our horizon is long enough. If the variable concerned has a trend – however small – then the forecasted variable will eventually become implausibly high or low, and any results based on such forecasts will lose their credibility. When forecasting variables in the long run, we must therefore rule out trends or impose arbitrary bounds on the variables being forecasted. However, if we impose arbitrary bounds, then the forecasting procedure becomes irrelevant, as we know the forecasted variable will eventually hit one of its bounds, and we may as well impose arbitrary values in the first place.

The implication is that we can only forecast our variables as they move around a zero trend, but in that case, why not just assume that the variable being forecasted takes its current value, or perhaps some typical recent value? Even if we had day-to-day forecasts, their fluctuations will tend to cancel out as the forecasted variable keeps returning toward its zero trend; a horizon-average of day-to-day forecasts would give us much the same result as projecting some recent value over our horizon period, and particularly so over longer horizons where the averaging-out process has more scope. Attempting to forecast these variables on a day-to-day basis is therefore pointless.[10]

---

[10] Forecasting volatility is also very difficult and, as Christoffersen *et al.* (1998, p. 109) conclude in a recent study, "Volatility forecastability seems to decline quickly with horizon, and seems to have largely vanished beyond horizons of ten or fifteen trading days." As the same study also points out, the temporal aggregation properties of existing volatility-forecasting models are not well understood, so

The foregoing discussion suggests that attempts to forecast the mean or volatility of returns over successive small periods are likely to be both difficult and unnecessarily complex. If forecasting with trends leads to explosive results over long horizons, and if fluctuations around a zero trend tend to cancel out, then we might as well use a simplistic approach and take a view about the average long-term values of the relevant parameters – which is exactly the approach adopted in the previous section.

5. Conclusions

This paper offers an easily implementable approach to the estimation of long-term VaR. This approach also provides some useful insights about the factors that determine long-term VaR and, in particular, about the impact of mean and volatility assumptions on estimates of long-term VaR. Our approach avoids problems associated with the square-root rule, as well as those associated with attempting to extrapolate day-to-day volatility forecasts over long horizons. Nonetheless, we should keep in mind that estimates of long-term VaR, like those of its short-term counterpart, are likely to be subject to considerable model and parameter risk.

---

we can rarely, if ever, rely on the alternative of temporal aggregation to obtain volatility forecasts.

References

Basle Committee on Banking Supervision (1996) *Amendment to the Capital Accord to Incorporate Market Risks*. Basle: Basle Committee on Banking Supervision, Bank for International Settlements, January.

Blake, D., A. Cairns, and K. Dowd (2000) "Extrapolating VaR by the Square-Root Rule." *Financial Engineering News*, August 2000, pp. 3, 7.

Blake, D., A. Cairns, and K. Dowd (2001) "PensionMetrics: Stochastic Pension Plan Design and Value-at-Risk during the Accumulation Phase." *Insurance: Mathematics and Economics* 29 (October): 187-215.

Christoffersen, P. F., F. X. Diebold, and T. Schuermann (1998) "Horizon Problems and Extreme Events in Financial Risk Management." Federal Reserve Bank of New York *Economic Policy Review*, October, pp. 109-118.

Dowd, K. (2002) *Measuring Market Risk*. Chichester and New York: Wiley and Sons.

Gupta, F., E. Stubbs, and Y. Thambiah (2000) "U.S. corporate pension plans: a value at risk analysis." *Journal of Portfolio Management* 26 (Summer): 65-72.

Panning, W. H. (1999) "The Strategic Uses of Value at Risk: Long-Term Capital Management for Property/Casualty Insurers." *North American Actuarial Journal*, Vol. 3, No. 2, pp. 84-105.