Journal of Banking & Finance 37 (2013) 847–853

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

On the role of the estimation error in prediction of expected shortfall

Carl Lönnbark *

Department of Economics, Umeå School of Business and Economics, Umeå University, Sweden

ARTICLE INFO

Article history:
Received 20 January 2012
Accepted 25 October 2012
Available online 8 November 2012

JEL classification:
G19
C52
C53
C58
G10

Keywords:
Backtesting
Delta method
Finance
GARCH
Risk management

ABSTRACT

In the estimation of risk measures such as Value at Risk and Expected shortfall relatively short estimation windows are typically used rendering the estimation error a possibly non-negligible component. In this paper we build upon previous results for the Value at Risk and discuss how the estimation error comes into play for the Expected Shortfall. We identify two important aspects where it may be of importance. On the one hand there is in the evaluation of predictors of the measure. On the other there is in the interpretation and communication of it. We illustrate magnitudes numerically and emphasize the practical importance of the latter aspect in an empirical application with stock market index data.

1. Introduction

The recent financial crisis highlights the need of properly understanding and measuring financial risks and in particular of evaluating the means of doing so. When it comes to measuring financial risk the Value-at-Risk (*VaR*) has during the past two decades or so emerged as the standard approach and it is today extensively employed by financial institutions over the world. This popularity is at least partially due to the fact that regulators have adopted the measure as a base for capital adequacy calculations. This was first stipulated in the 1996 Amendment to the first Basel Accord on banking supervision and later further detailed and reinforced in the second Accord (see Basel Committee on Banking Supervision, 2005; Basel Committee on Banking Supervision, 2006). In the aftermath of the financial crises new regulations have been developed to further strengthen capital requirement calculations (see Basel Committee on Banking Supervision, 2012b). Consequently, the measure has been given due attention in the literature (see Jorion (2007), for an extensive overview).

The *VaR* gives a potential portfolio loss that will only be exceeded with some (small) probability over a given horizon. As such it is conceptually simple. However, critique has been directed at the *VaR* measure both from the academia and from the industry. A complaint from the latter is that the *VaR* is silent about the size of the loss when it exceeds the *VaR*. Furthermore, the *VaR* may fail to acknowledge so-called tail risk. That is, two portfolios may have the same risk in terms of *VaR*, but their outcome in case of *VaR* exceedence may be substantially different (e.g. Yamai and Yoshiba, 2005). In an important paper Artzner et al. (1999) give a formal discussion of what constitutes a good risk measure and establish some properties of coherence that should be satisfied. In particular, a risk measure should acknowledge the principle of diversification. However, it is possible to find perverse cases, where the *VaR* does not satisfy this property. A measure that fares better in these respects is the Expected Shortfall (*ES*) that gives the expected loss given that the loss exceeds the *VaR*. As the name implies it says something about the size of the loss when disaster strikes, and it also acknowledges tail risk in a better way than *VaR*. The measure also possesses the desirable property of coherence. In fact, in a recent report the Basel Committee on Banking Supervision suggests a move towards the *ES* as the risk measure of choice for capital adequacy calculations (see Basel Committee on Banking Supervision, 2012a).

In computing the *VaR* and the *ES* a model for the joint movements of the risk factors of the portfolio is typically postulated and the parameters of that model are estimated based on a data-set containing past observations. Thus, uncertainty in the predictors of *VaR* and

* Tel.: +46 735 042030.
E-mail address: carl.lonnbark@econ.umu.se

0378-4266/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2012.10.013

ES arises from two primary sources. First of all, the true data generating process is not known, which gives rise to model risk. Secondly, the fact that the parameters of the hypothesized model must be estimated gives rise to estimation risk. Here, the focus is on the estimation risk. This source of error is often referred to as a second order issue and neglected though. Consequently, it is relatively understudied. In fact, Lan et al. (2007) report that the research on the uncertainty of *VaR* only amounts to about 2.5% of the *VaR* literature. In practise though, relatively short estimation windows of one or two years are typically used rendering the estimation error a non-negligible component. Indeed, the importance of estimation risk in this context has previously been emphasized by Jorion (1996) and Christoffersen and Gonçalves (2005) and others. In fact, Lönnbark (2010) demonstrates that the estimation error in *VaR* predictors may cause underestimation of portfolio risk in the sense that the probability of exceeding the estimated *VaR* is higher than the chosen level. Thus, the estimation error affects the interpretation of the *VaR*. In addition, when it comes to assessing the adequacy of a *VaR* model the conventional way is to compare a time series of historical *VaR* predictions to the corresponding portfolio returns. This procedure is commonly referred to as backtesting (e.g. Christoffersen, 2003, Ch. 8). A good *VaR* model should have a proportion of *VaR* exceedences (days when the loss exceeds the *VaR*) close to the chosen probability level. Consequently, as discussed in Escanciano and Olmo (2010) the estimation error also affects the backtesting procedure and may bias the breach frequency. Of obvious interest is what the picture looks like for the *ES* measure, which is the focus of this paper.

2. ES and VaR predictors

We assume that portfolio returns are generated in discrete time by

$$y_t = \mu(\boldsymbol{\theta}_{10}, I_{t-1}) + \sigma(\boldsymbol{\theta}_{20}, I_{t-1})\varepsilon_t, \tag{1}$$

where we take ε_t to be a standard normally distributed random variable. The $\mu(\cdot)$ and the $\sigma(\cdot)$ are the conditional mean and standard deviation functions, respectively. The vectors $\boldsymbol{\theta}_{10}$ and $\boldsymbol{\theta}_{20}$ contain true parameters and the set I_t contains the information available at time t. Typically, $\sigma(\cdot)$ is postulated indirectly in terms of the conditional variance (cf. the workhorse GARCH (1,1) specification of Bollerslev (1986) that parameterizes the conditional variance by $\sigma_t^2 = \beta_0 + \beta_1 y_{t-1}^2 + \beta_2\sigma_{t-1}^2$). For a portfolio with returns generated by (1) the one period ahead conditional *VaR*, VaR_t^α, satisfies $\Pr_{t-1}(y_t \leqslant -VaR_t^\alpha) = \alpha$, where the subscript $t-1$ indicates that the probability is conditional on I_{t-1}, and is in this case given explicitly by

$$VaR_t^\alpha = -\mu(\boldsymbol{\theta}_{10}, I_{t-1}) - \sigma(\boldsymbol{\theta}_{20}, I_{t-1})\Phi_\alpha^{-1}, \tag{2}$$

where Φ_α^{-1} is the inverse of the cdf of the standard normal distribution evaluated at α. The associated *ES* is given by

$$\begin{aligned} ES_t^\alpha &= -E_{t-1}(y_t | y_t \leqslant -VaR_t^\alpha) \\ &= -\mu(\boldsymbol{\theta}_{10}, I_{t-1}) - \sigma(\boldsymbol{\theta}_{20}, I_{t-1})\phi(\Phi_\alpha^{-1})/\alpha \end{aligned} \tag{3}$$

where $\phi(\cdot)$ is the pdf of the standard normal distribution and where the subscript $t-1$ on the expectation operator indicates that it is conditional on I_{t-1}. The *VaR* and the *ES* are conventionally reported as positive numbers. Hence, the minus signs in the definitions above.

When it comes to the estimation of the parameter vector, $\boldsymbol{\theta}_0 = (\boldsymbol{\theta}_{10}',\ \boldsymbol{\theta}_{20}')'$, the maximum likelihood estimator is commonly employed. It takes as the estimator the parameter vector, $\boldsymbol{\theta} = (\boldsymbol{\theta}_1',\ \boldsymbol{\theta}_2')'$, that maximizes the (conditional) likelihood function, $L \propto -(1/2)\sum(\ln\sigma_t^2 + (y_t - \mu_t)^2/\sigma_t^2)$, where $\mu_t = \mu\ (\boldsymbol{\theta}_1, I_{t-1})$ and $\sigma_t = \sigma\ (\boldsymbol{\theta}_2, I_{t-1})$. Given some regularity conditions the estimator vector, $\hat{\boldsymbol{\theta}}$, is asymptotically normally distributed with the true parameter vector, $\boldsymbol{\theta}_0$, as its mean and covariance matrix $\boldsymbol{\Sigma} = -[E(\partial^2 \ln L(\boldsymbol{\theta}_0)/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}')]^{-1}$. Predictors of VaR_t^α and ES_t^α are simply obtained by plugging in the estimator vector, $\hat{\boldsymbol{\theta}}$, in the expressions (2) and (3), respectively, to obtain

$$\widehat{VaR}_t^\alpha = -\mu(\hat{\boldsymbol{\theta}}_1, I_{t-1}) - \sigma(\hat{\boldsymbol{\theta}}_2, I_{t-1})\Phi_\alpha^{-1}, \tag{4}$$

and

$$\widehat{ES}_t^\alpha = -\mu(\hat{\boldsymbol{\theta}}_1, I_{t-1}) - \sigma(\hat{\boldsymbol{\theta}}_2, I_{t-1})\phi(\Phi_\alpha^{-1})/\alpha. \tag{5}$$

3. The role of the estimation error

When it comes to quantifying the uncertainty due to the estimation error in the *VaR* and the *ES* predictors we may rely on the asymptotic normality of the parameter estimator (cf. Hansen, 2006, and others) . Heuristically, asymptotic normality of $\widehat{VaR}_t^\alpha$ and $\widehat{ES}_t^\alpha$ follows from the asymptotic normality of $\hat{\boldsymbol{\theta}}$. We have

$$\widehat{VaR}_t^\alpha \sim N(VaR_t^\alpha, \delta_t^2), \tag{6}$$

and

$$\widehat{ES}_t^\alpha \sim N(ES_t^\alpha, \upsilon_t^2), \tag{7}$$

where the variances, δ_t^2 and υ_t^2, may be obtained by employing the delta method. In the sequel we maintain the assumption that $\widehat{VaR}_t^\alpha$ and $\widehat{ES}_t^\alpha$ are normally distributed. A key insight is that, in practice, we use a random predictor of the true *VaR* and when it comes to interpreting and communicating the measure the relevant probability is $\Pr_{t-1}\{y_t \leqslant -\widehat{VaR}_t^\alpha\}$. Clearly, this probability does not necessarily equal $\alpha$and may in fact equal some $\alpha^* > \alpha$ implying an underestimation of portfolio risk. Indeed, statements such as "the probability that the portfolio loss is less than the *VaR* is 100α% " may be quite misleading. In Fig. 1 we depict a situation with an unbiased *VaR* predictor.

For a *VaR* "draw" to the left of (minus) the true *VaR* the probability of exceedence is smaller than α. For a draw to the right the opposite is true. As the return density is positively sloped through the *VaR* density the latter will dominate. We note that if the return density were flat through the *VaR* density there would be no effect on the exceedence probability, i.e. $\alpha^* = \alpha$. Extrapolating on this reasoning we may conjecture that the difference between α^* and α is smaller for fat tailed return distributions.

Essentially, in the backtesting of a *VaR* predictor we compare draws from the *VaR* distribution to draws from the return distribution. Thus, the discussion above have a bearing on this procedure and for *VaR* the role of estimation error is essentially the same

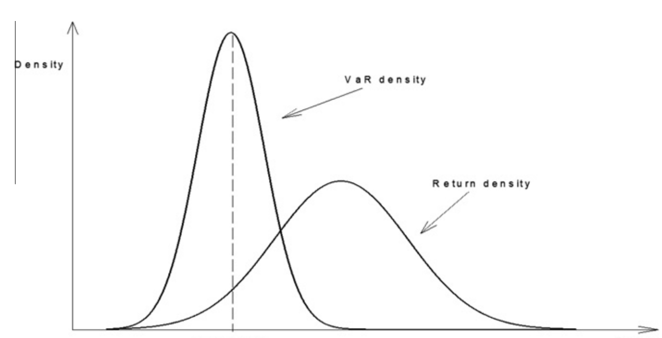

Fig. 1. *VaR* and return densities. *VaR* density and return density refers to the conditional densities of the *VaR* predictor and the return, respectively.