Journal of Banking & Finance 34 (2010) 2065–2076

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

Measuring portfolio credit risk correctly: Why parameter uncertainty matters

Nikola Tarashev *

Bank for International Settlements, Centralbahnplatz 2, 4002 Basel, Switzerland

A R T I C L E I N F O

Article history:
Received 23 June 2009
Accepted 26 January 2010
Available online 1 February 2010

JEL classification:
G20
G32
C11

Keywords:
Correlated defaults
Estimation error
Risk management

A B S T R A C T

Why should risk management systems account for parameter uncertainty? In addressing this question, the paper lets an investor in a credit portfolio face non-diversifiable uncertainty about two risk parameters – probability of default and asset-return correlation – and calibrates this uncertainty to a lower bound on estimation noise. In this context, a Bayesian inference procedure is essential for deriving and analyzing the main result, i.e. that parameter uncertainty raises substantially the tail risk perceived by the investor. Since a measure of tail risk that incorporates parameter uncertainty is computationally demanding, the paper also derives a closed-form approximation to such a measure.

1. Introduction

Measures of credit risk are often based on an analytic model and on the assumption that the parameters of this model are known with certainty. In turn, it is common practice for risk management systems to rely on such measures because of their tractability, even though it is attained by ignoring *estimation noise*. This practice may impair severely the quality of risk management systems because, besides credit-risk factors, estimation noise is another important determinant of uncertainty about potential losses.

In order to substantiate this claim, this paper generalizes the popular asymptotic single risk factor (ASRF) model of portfolio credit risk by allowing for noisy estimates of two key parameters: probability of default (PD) and asset-return correlation.[1] Applied to a stylized empirical framework, the generalized model delivers two alternative measures of tail risk that help underscore the importance of estimation noise. The first measure is a *naive* value-at-risk (VaR) of the portfolio, which accounts for the credit-risk factor but treats point estimates of the PD and asset-return correlation as equal to the true values of the respective parameters. The second is the *correct* VaR measure, which accounts not only for the credit-risk factor that influences the naive VaR but also for parameter uncertainty.

In principle, the correct VaR not only differs from the naive one but also does not match the *actual* risk that an investor is exposed to. The reason is that measuring the actual VaR of a portfolio requires knowledge of the true values of the risk parameters. By contrast, the correct VaR reflects the level of tail risk *perceived* by an investor, given his imperfect information about risk parameters.

Quantifying the correct VaR on the basis of a Bayesian inference procedure and comparing it to the naive one reveals that ignoring estimation noise should be expected to lead to a substantial understatement of the level of tail risk perceived by an investor. In the benchmark specification – where an investor in a homogeneous portfolio estimates the PD and asset-return correlation at 1% and 20%, respectively, on the basis of data covering 200 obligors over 10 years – the correct VaR is 27% higher than the corresponding naive VaR. This result is striking, not least because the underlying stylized empirical framework incorporates a lower bound on the amount of estimation noise.

In addition, accounting for estimation noise dampens (correctly) the sensitivity of VaR measures to changes in parameter estimates. The flip side is that, by abstracting from estimation noise, naive VaRs overstate the information content of parameter estimates. This manifests itself in that the difference between the correct and naive VaRs – i.e. the *add-on* induced by parameter uncertainty – decreases (increases) by less than the naive VaR when changes in parameter estimates suggest lower (higher) tail risk of the portfolio.

In comparison to Löffler (2003) and Gössl (2005) – which also incorporate estimation noise in measures of portfolio tail risk – this

* Tel.: +41 61 280 9213; fax: +41 61 280 9100.
E-mail address: nikola.tarashev@bis.org

[1] The paper abstracts from issues arising from *model* uncertainty. Kerkhof et al. (2010), Hsien-Hsing et al. (2009) and Tarashev and Zhu (2008a) study the impact of model mis-specification on measures of portfolio risk.

0378-4266/$ - see front matter
doi:10.1016/j.jbankfin.2010.01.017

paper conducts the analysis in a more transparent framework that allows for comparing in a straightforward fashion the importance of different sources of noise. Namely, a lengthening of the time series of the available data from 5 to 10 or from 10 to 20 years is seen to reduce the correct VaR add-on by a factor of two. In comparison, similar changes to the size of the cross section have a markedly smaller impact on the add-on. This finding is rooted in the standard assumption that credit risk is driven by asset returns that are serially uncorrelated but are correlated across obligors, which implies that increasing the time series of the data brings in more information than expanding the cross section.

The transparent framework of this paper also helps to analyze the trade-off between accuracy and reduction of the computational burden. This trade-off underscores the advantages of an approximate VaR measure, which exists in closed form and accounts for uncertainty about the PD but needs to be adjusted to reflect noise in observed asset returns. Given a judicious adjustment for such noise, this measure approximates the correct VaR quite well and alleviates substantially the computational burden.

That said, owing to the underlying Bayesian inference procedure, the computational burden is substantial for both the correct and the approximate VaR measures. This procedure is instrumental for capturing an important empirical regularity. Namely, over a realistic range of parameter values, higher levels of the PD and asset-return correlation are associated with greater noise in the associated estimates. This dependence between estimation noise and true parameter values, which is missed if one circumvents Bayesian inference, raises the probability that the true PD and correlation are bigger than their point estimates and, consequently, raises the correct VaR.

From a general point of view, this paper provides support to a call made by Borio and Tsatsaronis (2004) for including "measurement error information" in financial reporting. In deriving an ideal information set for attaining efficiency of the financial system, that paper emphasizes measurement error – a specific example of which is estimation error – as a piece of information that is of natural interest to risk managers and supervisors alike. From this perspective, the results derived below highlight specific scenarios in which knowledge of measurement error is indeed highly valuable as such error accounts for much of the uncertainty about potential credit losses.

The present paper differs in an important way from a number of recent articles – e.g. Tarashev and Zhu (2008a) and Heitfield (2008) – that have analyzed estimation noise in the context of portfolio credit risk.[2] Conditioning on hypothesized true parameter values, these articles study the extent to which ignoring estimation noise introduces errors in VaR measures.[3] However, the articles do not demonstrate how to *incorporate* the inevitable uncertainty about the true parameter values in measures of portfolio credit risk. In terms of the terminology introduced here, these articles quantify the discrepancy between the naive and actual VaRs but do not derive correct VaRs.

The rest of the paper is organized as follows: Section 2 describes the model and then derives alternative measures of portfolio VaR. These measures are considered in the context of an empirical framework that is outlined in Section 3. In turn, Section 4 presents and analyzes the quantitative results, paying particular attention to the trade-off between accuracy and computational complexity. Finally, Section 5 provides two extensions of the baseline analysis.

2. Stylized credit portfolio

The impact of parameter uncertainty on measures of tail risk is analyzed on the basis of a stylized credit portfolio. There are n exposures in this portfolio and all of them are of equal size, which is set to $1/n$. The analysis considers the limit $n \to \infty$, in which the portfolio is referred to as *asymptotic* or "perfectly fine-grained".[4]

The portfolio is also *homogeneous* in the sense that all of the exposures have the same credit characteristics. These characteristics are captured by two parameters: the degree to which a common risk factor affects an obligor's assets, denoted by ρ^*; and the obligor's probability of default, PD^*.[5] When parameter estimates are based on portfolios observed at different points in time (see Section 3.1), the true values of ρ^* and PD^* are assumed to remain constant across such portfolios.

The role of the two credit-risk parameters is rooted in the stochastic process followed by the value of the assets of each obligor, i:

$$\ln(V_{i,\tau}) = \ln(V_{i,\tau-\Delta}) + \mu^*\Delta + \sigma^*\sqrt{\Delta}\left(\sqrt{\rho^*}M_\tau + \sqrt{1-\rho^*}Z_{i,\tau}\right), \qquad (1)$$

$$\text{where } M_\tau \sim N(0,1), \quad Z_{i,\tau} \sim N(0,1),$$
$$Cov(M_\tau, Z_{i,\tau}) = 0, \quad Cov(Z_{i,\tau}, Z_{j,\tau}) = 0 \text{ for all } i \text{ and } j \neq i$$

and Δ denotes the period, in years, between two observations. There are two serially uncorrelated factors: one *common* to *all* obligors in the portfolio, M, and one specific to this obligor, Z_i. The drift of the asset value, the volatility of asset returns and the correlation between the asset returns of any two obligors – determined by μ^*, $\sigma^* > 0$ and $\rho^* \in [0,1]$, respectively – are the same for all i.

Obligor i defaults if and only if $\ln(V_{i,\tau})$ is below some threshold, D_i^*. In line with Merton (1974), default events are assumed to occur only at the end, $t \in \{1,2,\ldots,T\}$, of non-overlapping and adjacent *1-year* periods, which may be longer than the periods between two consecutive observations of the obligors' assets, i.e. $1 \geqslant \Delta$. Then, assuming that the loss-given-default on each exposure is unity, Eq. (1) implies that the loss on this portfolio over the next year t, is given by $L_{n,t}$:

$$L_{n,t} = \sum_i^n U_{i,t}, \text{ where} \qquad (2)$$
$$U_{i,t} = \begin{cases} 1/n & \text{if } \sqrt{\rho^*}M_t + \sqrt{1-\rho^*}Z_{i,t} < \Phi^{-1}(PD^*), \\ 0 & \text{otherwise,} \end{cases}$$
$$\Phi^{-1}(PD^*) = \left(D_i^* - \ln(V_{i,t-1}) - \mu^*\right)/\sigma^*$$

and PD^* is the unconditional 1-year probability of default, which is assumed to be the same across exposures (requiring that so is $D_i^* - \ln(V_{i,t-1})$). The expression $\Phi^{-1}(PD^*)$ is henceforth referred to as the (standardized) default point.[6]

An investor is interested in the maximum portfolio loss that is exceeded within a year with probability α, i.e. in the 1-year VaR at the $(1-\alpha)$ confidence level. It will be assumed that the investor

---

[2] Of the related articles, only Löffler (2003) analyzes uncertainty about PD on the basis of actual default data and does so via non-parametric bootstrap. See Lando and Skodeberg (2002) and Hanson and Schuermann (2006) for an extensive analysis of bootstrap approaches to the derivation of PD confidence intervals and Cantor et al. (2008) for an application of such an approach to a large dataset. The analysis below, just like Heitfield (2008) and Tarashev and Zhu (2008a), circumvents the use of bootstrap methods by assuming that the functional form, albeit not the parameter values, of the data-generating process is known.

[3] Bongaerts and Charlier (2009) perform a similar exercise in the context of regulatory capital assessments.

[4] See Gordy and Lütkebohmert (2007) for an analysis of real-life departures from perfect granularity.

[5] In this paper, "obligor" and "exposure" are used as close synonyms.

[6] In order to streamline the analysis, this setup abstracts from some important aspects of portfolio credit risk. In addition to adopting the afore mentioned assumption of perfect granularity, the setup rules out stochastic shocks to loss-given-default and exposure-at-default, which are modelled as correlated with the common default-risk factor by Kupiec (2008). The setup also abstracts from cross-sectional dispersion of credit-risk parameters, which is analyzed by Tarashev and Zhu (2008a).