Dependent default and recovery: Markov chain Monte Carlo study of downturn Loss Given Default credit risk model

Pavel V. Shevchenko[1]   Xiaolin Luo[2]

(Received 6 January 2012; revised 22 May 2012)

Abstract

There is empirical evidence that recovery rates tend to go down just when the number of defaults goes up in economic downturns. This has to be taken into account in estimation of the capital against credit risk required by Basel II to cover losses during the adverse economic downturns; the so-called "downturn Loss Given Default" requirement. This article presents a methodology for estimation of the Loss Given Default credit risk model with the default and recovery dependent via the latent systematic risk factor using a Bayesian inference approach and Markov chain Monte Carlo method. This approach allows joint estimation of all model parameters and latent systematic factor, and all relevant uncertainties. For illustration, we fit the model using Moody's annual default and recovery rates for corporate bonds for the period 1982–2010.

Contents

1 Introduction

Default and recovery rates are key components of Loss Given Default (LGD) models proposed for calculation of economical capital (EC) against credit risk. The classic LGD model implicitly assumes that the default rates and recovery rates are independent. Motivated by empirical evidence that recovery rates tend to go down just when the number of defaults goes up in economic downturns, Frye [3], Pykhtin [9] and Düllmann and Trapp [2] extended the classic model to include dependence between the default and recovery via a common systematic factor. These models have been suggested by some

banks for assessment of the Basel II "downturn LGD" requirement [1]. The Basel II "downturn LGD" reasoning is that recovery rates may be lower during economic downturns when default rates are high; and that a capital should be sufficient to cover losses during these adverse circumstances. The extended models represent an important enhancement of credit risk models used in earlier practice, such as *CreditMetrics* and *CreditRisk+*, that do not account for dependence between default and recovery.

Publicly available data provided by Moody's or Standard&Poor's rating agencies are annual averages of defaults and recoveries. These data are of limited size, covering a couple of decades at most. As shown in this article, the impact of the parameter uncertainty on capital estimate can be very significant, due to the limited data size. To our knowledge, the quantitative impact of parameter uncertainty has not been specifically addressed. Increasingly, quantification of parameter uncertainty and its impact on EC has become a key component of financial risk modeling and management; Luo et al. [5] and Peters et al. [8] recently gave examples in operational risk and insurance. This article studies parameter uncertainty and its impact on EC estimate in the LGD model, where default and recovery are dependent via the latent systematic risk factor. We demonstrate how the model can be estimated using the Bayesian approach and Markov chain Monte Carlo (MCMC) method. This approach allows joint estimation of all model parameters and latent systematic factor, and all relevant uncertainties.

2   Loss Given Default model

Following previous studies [2, 3, 9], consider a homogenous portfolio of J borrowers over a chosen time horizon. To avoid cumbersome notation, we assume that the jth borrower has one loan with principal amount A_j. The loss rate (loss amount relative to the loan amount) of the portfolio due to

defaults is

$$L = \sum_{j=1}^{J} w_j L_j = \sum_{j=1}^{J} w_j I_j \max(1 - R_j, 0), \tag{1}$$

where w_j is the weight of loan j, $w_j = A_j / \sum_{m=1}^{J} A_m$; L_j is the loss rate of loan j due to potential default; $1 - \max(1 - R_j, 0)$ is the recovery rate of loan j after default; I_j is an indicator variable associated with the default of loan j, $I_j = 1$ if firm j defaults, otherwise $I_j = 0$. In general R_j is not the same as the recovery rate since the latter is subject to a cap of one.

Denote the probability of default for firm j by p, that is, $\Pr[I_j = 1] = p$. Let C_j be an underlying latent random variable (financial well-being) such that firm j defaults if $C_j < \Phi^{-1}(p)$, where $\Phi(\cdot)$ is the standard normal distribution and $\Phi^{-1}(\cdot)$ is its inverse. That is, $I_j = 1$ if $C_j < \Phi^{-1}(p)$ and $I_j = 0$ otherwise. The value C_j for each firm depends on a systematic risk factor X and a firm specific (idiosyncratic) risk factor Z_j^C as

$$C_j = \sqrt{\rho} X + \sqrt{1 - \rho} Z_j^C, \tag{2}$$

where $Z_1^C, \ldots, Z_J^C$ are all independent. Also, X and Z_j^C are assumed independent and from the standard normal distribution. Conditional on X, the financial conditions of any two firms are independent. Unconditionally, ρ is the correlation between the financial conditions of two firms.

Several studies [2, 3, 9] considered normal, lognormal and logit-normal distributions for the recovery. It was shown by Düllmann and Trapp [2] that EC estimates from these three recovery models are very close to each other; the difference is within 2%. In addition, statistical tests favoured the normal distribution model. Thus we model the recovery rate as

$$R_j = \mu + \sigma\sqrt{\omega} X + \sigma\sqrt{1 - \omega} Z_j, \quad \omega \in [0, 1], \tag{3}$$

where X and Z_j are assumed independent and from the standard normal distribution. Also, Z_j and Z_j^C are assumed independent. The recovery and default processes are dependent via systematic factor X.

Algorithm 1: Quantile for given parameters.

1. Draw an independent sample from $\Phi(\cdot)$ for the systematic factor X.
2. For each j, draw Z_j^C from $\Phi(\cdot)$; calculate C_j and I_j.
3. For each j, draw Z_j from $\Phi(\cdot)$, find $R_j = \mu + \sigma\sqrt{\omega}X + \sigma\sqrt{1-\omega}Z_j$.
4. Find loss L for the entire portfolio using (1), that is, a sample from $F_L(\cdot \mid \theta)$.
5. Repeat steps 1–4 to obtain N samples of L.
6. Estimate $Q_q(\theta)$ using obtained samples of L in the standard way.

---

3   Economic capital

It is common to define the EC for credit risk as a high quantile of the distribution of loss L, that is,

$$Q_q(\theta) \equiv Q_q = \inf\{z : \Pr[L > z \mid \theta] \leqslant 1 - q\} = \inf\{z : F_L(z \mid \theta) \geqslant q\}, \quad (4)$$

where q is a quantile level; $F_L(z \mid \theta)$ is the distribution function of the loss L with the density denoted as $f_L(z \mid \theta)$; and $\theta = (p, \rho, \mu, \sigma, \omega)$ is the vector of model parameters.

The EC measured by the quantile $Q_q(\theta)$ is a function of θ. Typically, given observations, the maximum likelihood estimator (MLE) $\hat{\theta}$ is used as point estimate for θ. Then, the loss density for the next time period is estimated as $f_L(z \mid \hat{\theta})$ and its quantile, $Q_q(\hat{\theta})$, is used for EC calculation. The distribution of L is not tractable in closed form for an arbitrary portfolio. In this case the Monte Carlo method can be used with the logical steps specified in Algorithm 1.

Bank loans are subject to the borrower specific risk and systematic risk. In the case of a diversified portfolio with a large number of borrowers, the idiosyncratic risk can be eliminated and the loss depends on X only. Gordy [4] showed that the distribution of portfolio loss L has a limiting form as $J \to \infty$, provided that each weight w_j goes to zero faster than $1/\sqrt{J}$. The limiting

loss rate L^∞ is given by the expected loss rate conditional on X:

$$L^\infty \equiv L^\infty(X) = \sum_{j=1}^{J} w_j E[I_j \mid X] E[\max(1 - R_j, 0) \mid X] = \Lambda(X) S(X), \quad (5)$$

where $\Lambda(X) = E[I_j \mid X]$ is the conditional probability of default of firm j and $S(X) = E[\max(1 - R_j, 0) \mid X]$ is the conditional expected value of loss rate. That is, the distribution of L^∞ is fully implied by the distribution of X. Because $L^\infty(X)$ is a monotonic decreasing function and X is from the standard normal distribution, the quantile of $L^\infty(X)$ at level q can be calculated as $Q_q^\infty = L^\infty \left(X = \Phi^{-1}(1 - q) \right)$. Like Düllmann and Trapp [2], we define EC of the diversified portfolio loss distribution $L^\infty(X)$ as the 0.999 quantile

$$\mathrm{EC}^\infty = Q_{0.999}^\infty = L^\infty \left(\Phi^{-1}(0.001) \right) = \mathrm{PD} \times \mathrm{LGD}, \quad (6)$$

where $\mathrm{PD} = \Lambda(\Phi^{-1}(0.001))$ and $\mathrm{LGD} = S(\Phi^{-1}(0.001))$ are stressed probability of default (*stressed* PD) and stressed loss given default (*stressed* LGD) respectively. Using (2), the conditional probability of default is

$$\Lambda(X) = \Phi \left(\frac{\Phi^{-1}(p) - \sqrt{\rho} X}{\sqrt{1 - \rho}} \right). \quad (7)$$

Also, the expected conditional loss rate for the normally distributed recovery rate model (3) is easily calculated as

$$S(X) = \int_{-\infty}^{\infty} \max \left(1 - \mu - \sigma\sqrt{\omega} X - \sigma\sqrt{1 - \omega} z, 0 \right) f_N(z) \, dz$$

$$= \left(1 - \mu - \sigma\sqrt{\omega} X \right) \Phi(z_c) + \frac{\sigma\sqrt{1 - \omega}}{\sqrt{2\pi}} e^{-z_c^2/2}, \quad (8)$$

where $z_c = \left(1 - \mu - \sigma\sqrt{\omega} X \right) / \left(\sigma\sqrt{1 - \omega} \right)$ and $f_N(z)$ is the standard normal density. For the real data used in this study, it is well approximated as $S(X) \approx E[(1 - R_j) \mid X] = 1 - \mu - \sigma\sqrt{\omega} X$.

4 Likelihood function

Consider time periods $t = 1, \ldots, T$ (so that $T + 1$ corresponds to the next future year), where the following data of default and recovery for a loan portfolio of J_t firms are observed: D_t is the number of defaults in year t, and its realization is d_t; $\Psi_t = D_t/J_t$ is the default rate in year t, and its realization is ψ_t; $\overline{R}_t = \sum_{j=1}^{D_t} R_j(t)/D_t$ is the average recovery rate in year t, where $R_1(t), \ldots, R_{D_t}(t)$ are individual recoveries, and its realization is $\bar{r}_t$. Also, the systematic factor X corresponding to the time periods is denoted as $X_1, \ldots, X_{T+1}$ and its realization is $x_1, \ldots, x_{T+1}$. We assume that $X_1, \ldots, X_{T+1}$ are independent and all idiosyncratic factors (Z_j, Z_j^C) corresponding to the time periods are independent.

4.1 Exact likelihood function

The joint density of the number of defaults and average recovery rate $(D_t, \overline{R}_t)$ are calculated by integrating out the latent variable X_t for each t as

$$f(d_t, \bar{r}_t) = \int f(\bar{r}_t \mid d_t, x_t) f(d_t \mid x_t) f_N(x_t) dx_t, \tag{9}$$

where the conditional densities $f(d_t \mid x_t)$ and $f(\bar{r}_t \mid d_t, x_t)$ are derived as follows.

Given $X_t = x_t$, all firms in a homogenous portfolio have the same conditional default probability $\Pr[I_j(t) = 1 \mid X_t = x_t] = \Lambda(x_t)$ evaluated in (7). Thus, the conditional distribution of $D_t = \sum_{j=1}^{J_t} I_j(t)$ is binomial

$$f(d_t \mid x_t) = \Pr[D_t = d_t \mid X_t = x_t] = \binom{J_t}{d_t} [\Lambda(x_t)]^{d_t} [1 - \Lambda(x_t)]^{J_t - d_t}. \tag{10}$$

Often it is well approximated by the normal distribution $N(\mu_t, \sigma_t^2)$ with mean $\mu_t = J_t \Lambda(x_t)$ and variance $\sigma_t^2 = J_t \Lambda(x_t)(1 - \Lambda(x_t))$.

Conditional on $X_t = x_t$ and $D_t = d_t$, individual recoveries $R_1(t), \ldots, R_{d_t}(t)$ are independent and from $N(\mu_r, \sigma_r^2)$ with $\mu_r = \mu + \sigma\sqrt{\omega}\,x_t$ and $\sigma_r = \sigma\sqrt{1-\omega}$. Thus the average $\overline{R}_t$ is from $N(\mu_R, \sigma_R^2)$ with $\mu_R = \mu_r$ and $\sigma_R^2 = \sigma_r^2/d_t$, that is,

$$f(\overline{r}_t \mid d_t, x_t) = \frac{1}{\sqrt{2\pi}\sigma_R} \exp\left(-\frac{(\overline{r}_t - \mu_R)^2}{2\sigma_R^2}\right). \tag{11}$$

If recovery distribution is different from normal, then the average $\overline{R}_t$ can still be approximated by a normal distribution if d_t is large (and variance is finite). Define the data vectors $\mathbf{D} = (D_1, \ldots, D_T)$ and $\overline{\mathbf{R}} = (\overline{R}_1, \ldots, \overline{R}_T)$, then the joint likelihood function for data $\mathbf{D}$ and $\overline{\mathbf{R}}$ is

$$\ell_{\mathbf{D}, \overline{\mathbf{R}}}(\boldsymbol{\theta}) = \prod_{t=1}^{T} f(d_t, \overline{r}_t). \tag{12}$$

This joint likelihood function could estimate parameters $\boldsymbol{\theta}$ by MLE maximizing this likelihood. However, the likelihood involves numerical integration with respect to the latent variables $\mathbf{X}$. It is difficult to accurately compute these integrals, especially if the likelihood is used within numerical maximization procedures. A straightforward and problem-free alternative is to take the Bayesian approach and treat $\mathbf{X}$ in the same way as other parameters, and formulate the problem in terms of the likelihood conditional on $\boldsymbol{\gamma} = (\boldsymbol{\theta}, \mathbf{X})$. Then the required conditional likelihood is easily calculated as

$$\ell_{\mathbf{D}, \overline{\mathbf{R}}}(\boldsymbol{\gamma}) = \prod_{t=1}^{T} f(d_t \mid x_t, \boldsymbol{\theta}) f(\overline{r}_t \mid d_t, x_t, \boldsymbol{\theta}), \tag{13}$$

avoiding integration with respect to $\mathbf{X}$. Estimation based on this likelihood is discussed in detail in Section 5.

4.2   Approximate likelihood and closed form Maximum Likelihood Estimation

Assuming a large number of firms in the portfolio, some approximation is justified to find MLE for the likelihood (12). We adopt an approach

from Düllmann and Trapp [2], estimating the default process parameters $\theta_D = (\rho, p)$ and systematic factor X first, and then fitting the recovery parameters $\theta_R = (\mu, \sigma, \omega)$.

Given X_t, the conditional default probability $\Lambda_t = \Lambda(X_t)$ is a monotonic function of X_t; see (7). The density of X_t is the standard normal, thus the change of probability measure gives the density for Λ_t at $\Lambda_t = \lambda_t$:

$$f(\lambda_t \mid \theta_D) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x_t^2}{2}\right) \left|\frac{dx_t}{d\lambda_t}\right|, \tag{14}$$

where x_t is the function of λ_t, the inverse of (7),

$$x_t = \left(\Phi^{-1}(p) - \sqrt{1-\rho}\,\Phi^{-1}(\lambda_t)\right)/\sqrt{\rho}. \tag{15}$$

For year t we observe the default rate Ψ_t that for $J_t \to \infty$ approaches Λ_t. Therefore, the likelihood for observed default rates $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_T)$ is

$$\ell_D(\theta_D) = \prod_{t=1}^{T} f(\lambda_t = \psi_t \mid \theta_D). \tag{16}$$

Maximizing (16) gives the following MLE for ρ and p:

$$\hat{\rho} = \frac{\sigma_\delta^2}{1 + \sigma_\delta^2}, \quad \hat{p} = \Phi\left(\frac{\bar{\delta}}{\sqrt{1 + \sigma_\delta^2}}\right), \tag{17}$$

where $\bar{\delta} = \sum_{t=1}^{T} \delta_t/T$, $\sigma_\delta^2 = \sum_{t=1}^{T}(\delta_t - \bar{\delta})^2/T$ and $\delta_t = \Phi^{-1}(\psi_t)$. The factor X_t is then estimated using (15) with default parameters (p, ρ) replaced by MLE as

$$\hat{x}_t = \left(\Phi^{-1}(\hat{p}) - \sqrt{1-\hat{\rho}}\,\delta_t\right)/\sqrt{\hat{\rho}}. \tag{18}$$

Given X_t and D_t, the average recovery rate $\bar{R}_t$ is from $N(\mu_R, \sigma_R^2)$ with mean $\mu_R = \mu + \sigma\sqrt{\omega}X_t$ and variance $\sigma_R^2 = \sigma^2(1-\omega)/d_t$. Thus the likelihood for T observations of the average recovery rate $\bar{\mathbf{r}} = (\bar{r}_1, \ldots, \bar{r}_T)$ is

$$\ell_{\bar{R}}(\theta_R, \mathbf{x}) = \prod_{t=1}^{T} \sqrt{\frac{d_t}{2\pi\sigma^2(1-\omega)}} \exp\left(-\frac{d_t(\bar{r}_t - \mu - \sigma\sqrt{\omega}x_t)^2}{2\sigma^2(1-\omega)}\right). \tag{19}$$

Düllmann and Trapp [2] estimate θ_R by MLE via maximization of (19) with respect to θ_R, where x_t is replaced with $\hat{x}_t$. Due to numerical difficulties with the maximization, they estimated σ by the historical volatility $\hat{\sigma}_h$ of the recovery rate $\bar{r}_t$. However, re-parameterizing with $\sigma_1 = \sigma\sqrt{\omega}$ and $\sigma_2 = \sigma\sqrt{1-\omega}$, we derive the following closed form solutions for the MLE of (μ, σ, ω):

$$\hat{\sigma}_1 = \frac{\left(\sum_t d_t \bar{r}_t X_t\right)\left(\sum_t d_t\right) - \left(\sum_t d_t \bar{r}_t\right)\left(\sum_t d_t X_t\right)}{\left(\sum_t d_t X_t^2\right)\left(\sum_t d_t\right) - \left(\sum_t d_t X_t\right)^2}, \tag{20}$$

$$\hat{\mu} = \frac{\left(\sum_t d_t \bar{r}_t X_t\right) - \left(\sum_t d_t X_t^2\right)\hat{\sigma}_1}{\sum_t d_t X_t}, \quad \hat{\sigma}_2 = \sqrt{\frac{1}{T}\sum_t d_t (r_t - \hat{\mu} - \hat{\sigma}_1 X_t)^2}, \tag{21}$$

$$\hat{\omega} = \frac{\hat{\sigma}_1^2}{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}, \quad \hat{\sigma} = \sqrt{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}. \tag{22}$$

Note that numerical instabilities in the numerical maximization procedure appear because the variance of the average recoveries goes to zero for large number of defaults causing flatness of the likelihood. This problem would not appear if we fit individual recoveries; unfortunately data for individual recoveries are not available.

5 Bayesian inference and Markov chain Monte Carlo

The parameters θ are unknown and it is important to account for this uncertainty when the capital is estimated. A standard frequentist approach to estimate this uncertainty is based on limiting results of normally distributed MLE for large datasets. We take the Bayesian approach because the dataset is small and the parameter uncertainty distribution is very different from normal. From a Bayesian perspective, both the parameters θ and the latent factor $\mathbf{X}$ are random variables. Given a prior density $\pi(\gamma)$ and a data likelihood

$\pi(\mathbf{y} \mid \boldsymbol{\gamma}) = \ell_{\mathbf{Y}}(\boldsymbol{\gamma})$, where $\boldsymbol{\gamma} = (\boldsymbol{\theta}, \mathbf{X})$ and $\mathbf{Y} = (\mathbf{D}, \overline{\mathbf{R}})$ is the data vector, the density of $\boldsymbol{\gamma}$ conditional on $\mathbf{Y} = \mathbf{y}$ (posterior density) is determined by Bayes theorem as

$$\pi(\boldsymbol{\gamma} \mid \mathbf{y}) \propto \pi(\mathbf{y} \mid \boldsymbol{\gamma})\pi(\boldsymbol{\gamma}). \qquad (23)$$

The posterior is then be used for predictive inference and analysis of the uncertainties. There are many useful texts on Bayesian inference [10, e.g.], and there are many recent examples in operational risk and insurance [14, 8, 7].

The explicit evaluation of posterior (23) is often difficult and one can use the MCMC method to sample from the posterior. In particular, MCMC allows to get samples of $\boldsymbol{\theta}$ and $\mathbf{X}$ from the joint posterior $\pi(\boldsymbol{\theta}, \mathbf{X} \mid \mathbf{y})$. Then taking samples of $\boldsymbol{\theta}$ marginally, we get the posterior for model parameters $\pi(\boldsymbol{\theta} \mid \mathbf{y})$; that is, effectively integrating out the latent factor $\mathbf{X}$. Similarly, taking samples of $\mathbf{X}$ marginally, we get the posterior for systematic factor $\pi(X_t \mid \mathbf{y})$. The posterior mean is a commonly used point estimate. We adopt the component-wise Metropolis–Hastings algorithm for sampling from the posterior $\pi(\boldsymbol{\gamma} \mid \mathbf{y})$, following an established procedure [15, 8]. Other MCMC methods such as the univariate slice sampler [7] can also be used. For numerical efficiency, we work with parameter $\Phi^{-1}(p)$. Also, we assume a uniform prior for all parameters and the standard normal distribution as the prior for $X_1, \ldots, X_T$. The only subjective judgement we bring to the prior is the lower and upper bounds of the parameter values

$$\Phi^{-1}(p) \in (-10, 10), \quad \rho \in (0, 1), \quad \mu \in (0, 1), \quad \sigma \in (0.01, 1.0), \quad \omega \in (0, 1).$$

The parameter support range should be sufficiently large so that the posterior is implied mainly by the observed data. We checked that an increase in parameter bounds did not lead to any material difference in results.

The starting value of the chain for the kth component is set to a uniform random number drawn independently from the support (a_k, b_k). In the single component Metropolis–Hastings algorithm, we adopt a Gaussian density (truncated below a_k and above b_k) for the proposal density. For each component the variance parameter of the proposal was pre-tuned and adjusted

so that the acceptance rate is close to 0.234 (the optimal acceptance rate for d-dimensional target distributions with independent and identically distributed components [11]). The chain is run for $100,000$ posterior samples (after $20,000$ "burn-in" samples).

6 Bayesian capital estimates

Bayesian methods are particularly convenient to quantify parameter uncertainty and its impact on capital estimate [13]. Under the Bayesian approach, the full predictive density (accounting for parameter uncertainty) of the next time period loss L_{T+1}, given data $\mathbf{Y} = \mathbf{y}$, is

$$f_{L_{T+1}}(z \mid \mathbf{y}) = \int f_{L_{T+1}}(z \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \mathbf{y}) \, d\boldsymbol{\theta}, \tag{24}$$

assuming that L_{T+1} and $\mathbf{Y}$ are independent, given $\boldsymbol{\theta}$. Its quantile,

$$Q_q^P = \inf\{z : \Pr[L_{T+1} > z \mid \mathbf{Y}] \leqslant 1 - q\}, \tag{25}$$

can be used as a risk measure for EC. The procedure for simulating L_{T+1} from (24) and calculating Q_q^P is simple:

1. draw a sample of $\boldsymbol{\theta}$ from the posterior $\pi(\boldsymbol{\theta} \mid \mathbf{y})$, for example, using MCMC;

2. given $\boldsymbol{\theta}$, simulate loss L following steps 1–4 in Algorithm 1;

3. repeat steps 1–2 to obtain N samples of L;

4. estimate Q_q^P using samples of L in the standard way.

Another approach under a Bayesian framework to account for parameter uncertainty is to consider a quantile $Q_q(\boldsymbol{\theta})$ of the loss density $f_{L_{T+1}}(\cdot \mid \boldsymbol{\theta})$,

$$Q_q(\boldsymbol{\theta}) = \inf\{z : \Pr[L_{T+1} > z \mid \boldsymbol{\theta}] \leqslant 1 - q\}. \tag{26}$$

Given that $\boldsymbol{\theta}$ is distributed as $\pi(\boldsymbol{\theta} \mid \mathbf{y})$, one can find the associated distribution of $Q_q(\boldsymbol{\theta})$, form a predictive interval to contain the true quantile value with some probability and argue that the conservative estimate of the capital accounting for parameter uncertainty should be based on the upper bound of the interval. However, it might be difficult to justify the choice of the confidence level for the interval. The procedure to obtain the posterior distribution of quantile $Q_q(\boldsymbol{\theta})$ is simple:

1. draw a sample of $\boldsymbol{\theta}$ from the posterior $\pi(\boldsymbol{\theta} \mid \mathbf{y})$, for example, using MCMC;

2. compute $Q_q = Q_q(\boldsymbol{\theta})$ using, for example, Algorithm 1;

3. repeat steps 1–2 to obtain N samples of $Q_q(\boldsymbol{\theta})$.

For the limiting case of a large number of borrowers, Step 2 is approximated by a closed form formula.

The extra loading for EC due to parameter uncertainty is formally defined as the difference between the quantile of the full predictive distribution accounting for parameter uncertainty $Q_{0.999}^P$ and posterior mean of $Q_{0.999}(\boldsymbol{\theta})$, that is, $Q_{0.999}^P - E[Q_{0.999}(\boldsymbol{\theta})]$.

7   Numerical results using Moody's data

Using historical data for the overall corporate default and recovery rates over 1982–2010 from Moody's report [6], we fit the model using MCMC and MLE. Table 1 shows posterior summary and MLE for the model parameters (the coefficient of variation, CV, is defined as the ratio of standard deviation to the mean). Significant kurtosis and positive skewness in most parameters indicate that Gaussian approximation for parameter uncertainties is not appropriate. Also, all MLEs are within one standard deviation from the posterior mean. The posterior mean of the systematic factor X_t for year 2009 is about -2.27, which corresponds to approximately the 99% quantile level

of the diversified portfolio. This maximum negative systematic factor for year 2009 is a consequence of a disastrous year 2008 when the bankruptcy of Lehman Brothers occurred. Comparison of MLE and the posterior mean for the latent factor $\mathbf{X}$ is shown in Figure 1.

The MCMC predictions on stressed PD, LGD and EC in comparison with the corresponding MLE are shown in Table 2. The MLE for EC is 35% lower than the posterior mean, 24% lower than the posterior median, and more than 50% lower than the 0.75 quantile of the posterior for EC. The uncertainty in the posterior of EC is large, the CV is about 34.5%; also note a large difference between the 0.75 quantiles and the 0.25 quantiles of the EC posterior. Underestimation of the EC by MLE in comparison with posterior estimates is significant due to large parameter uncertainty and large skewness in the EC posterior. Also, we get the following results for the 0.999 quantile $Q^P_{0.999}$ of the full predictive loss density for portfolios with different number of borrowers J: $Q^P_{0.999} = (0.1454, 0.1092, 0.1026, 0.1026)$ for $J = (50, 500, 5000, \infty)$ respectively. The diversification effect when J increases is evident. In particular, $Q^P_{0.999}$ at $J = 500$ is about 25% lower than the case at $J = 50$; and for $J = 5000$ is virtually the same as for the limiting case $J = \infty$. Note that $Q^P_{0.999}$ at $J = \infty$ is about 50% larger than MLE for the EC, and about 15% larger than the posterior mean of $Q^\infty_{0.999}(\boldsymbol{\theta})$. The 15% impact of parameter uncertainty on the EC indicates that the 1982–2010 dataset is long enough for a more or less confident use of the model for capital quantification. Of course, these results should be viewed as an illustration of the estimation methodology rather than realistic estimates of the capital. The latter would require fitting similar data for the actual loan portfolio in a bank.

8   Conclusion

The presented methodology allows joint estimation of the model parameters and latent systematic risk factor in the well known LGD model via a Bayesian approach and the MCMC method. This approach allows an easy calculation

Table 1: MLE and MCMC posterior statistics of the model parameters.

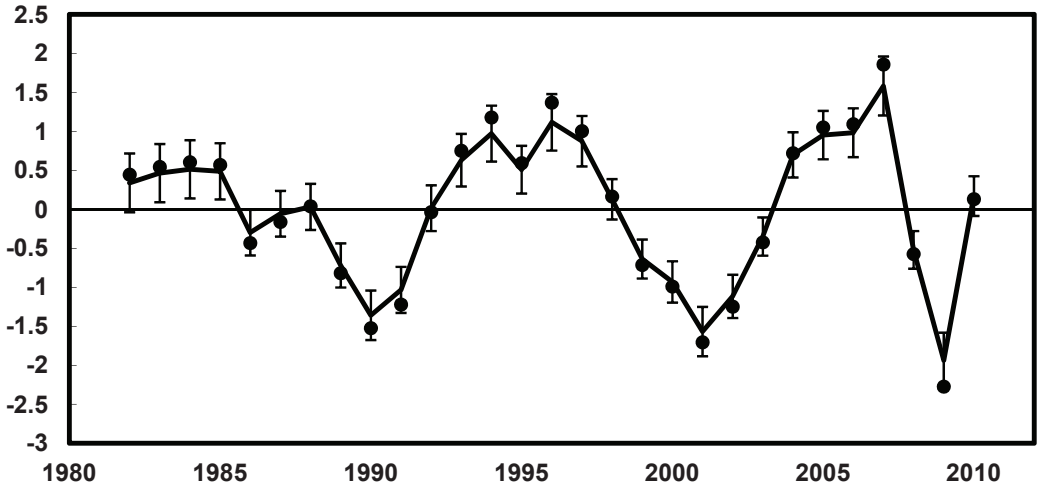| item	MLE	mode	mean	stdev	skewness	kurtosis	CV
p	0.0167	0.0177	0.0179	0.0028	0.812	4.62	0.154
ρ	0.0635	0.141	0.0815	0.024	1.01	4.35	0.286
μ	0.411	0.439	0.414	0.022	0.309	3.19	0.055
ω	0.0192	0.0717	0.031	0.016	1.24	5.39	0.51
σ	0.499	0.449	0.502	0.070	0.588	3.63	0.140

Figure 1: MLE (dots) and posterior mean (solid line) of systematic factor X_t. Error bars correspond to posterior standard deviation of X_t.

Table 2: MLE and MCMC posterior statistics for PD, LGD and EC.

item	MLE	mean	stdev	0.25Q	0.5Q	0.75Q	CV
PD	0.0819	0.103	0.029	0.0825	0.0968	0.116	0.288
LGD	0.803	0.847	0.0562	0.808	0.841	0.880	0.066
EC	0.0657	0.0888	0.031	0.0672	0.0814	0.101	0.345

of the full predictive loss density $f_{L_{T+1}}(\cdot \mid \mathbf{y})$ accounting for parameter uncertainty. Then the economic capital can be based on the high quantile of this distribution. For realistic estimates of economic capital, one should fit the model using dataset of defaults and recoveries for the actual loan portfolios in a bank. In this article, for illustration, we fitted the model to Moody's data for corporate bonds. The Moody's datasets of the annual defaults and recovery rates are small. Thus the parameter uncertainty is large and the posterior is very different from the normal distribution indicating that the Gaussian approximation for parameter uncertainties (typically used under the frequentist maximum likelihood approach assuming a large sample limit) is not appropriate. For simplicity, we assumed a homogeneous portfolio. The approach can be extended to deal with non-homogeneous portfolios, more than one latent factor, and mean reversion in the systematic factor. It should not be difficult to incorporate macroeconomic factors as done by Rösch and Scheule [12].

References

[1] Basel Committee on Banking Supervision. *Guidance on Paragraph 468 of the Framework Document.* Bank for International Settlements, Basel, July 2005. http://www.bis.org/publ/bcbs115.htm C187

[2] K. Düllmann and M. Trapp. Systematic risk in recovery rates - an empirical analysis of us corporate credit exposures. *Discussion Paper, Series2: Banking and Financial Supervision*, pages 1–44, 2004. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=494462 C186, C187, C188, C190, C193, C194

[3] J. Frye. Depressing recoveries. *Risk*, 13:106–111, 2000. http://www.risk.net/data/Pay_per_view/risk/technical/2000/risk_1100_creditrisk.pdf C186, C187, C188

[4] M. Gordy. A risk-factor foundation for ratings-based bank capital rules. *Finance and Economics Discussion Series 2002-55, Washington: Board of Governors of the Federal Reserve System*, 2002. doi:10.2139/ssrn.361302 C189

[5] X. Luo, P. V. Shevchenko, and J. Donnelly. Addressing impact of truncation and parameter uncertainty on operational risk estimates. *The Journal of Operational Risk*, 2(4):3–26, 2007. http://www.risk.net/journal-of-operational-risk/journal/ 2160945/journal-operational-risk-volume-number-winter-2007 C187

[6] Moody's. Corporate default and recovery rates, 1920–2010. Technical report, February 2011. http://www.moodys.com C197

[7] G. W. Peters, P. V. Shevchenko, and M. V. Wüthrich. Dynamic operational risk: modelling dependence and combining different data sources of information. *The Journal of Operational Risk*, 4(2):69–104, 2009. http://www.risk.net/journal-of-operational-risk/journal/ 2160936/journal-operational-risk-volume-number-summer-2009 C195

[8] G. W. Peters, P. V. Shevchenko, and M. V. Wüthrich. Model uncertainty in claims reserving within Tweedie's compound poisson models. *ASTIN Bulletin*, 39(1):1–33, 2009. doi:10.2143/AST.39.1.2038054 C187, C195

[9] M. Pykhtin. Unexpected recovery risk. *Risk*, 16(8):74–78, 2003. http://www.risk.net/risk-magazine/technical-paper/1530263/ unexpected-recovery-risk C186, C187, C188

[10] C. P. Robert. *The Bayesian Choice*. Springer Verlag, New York, 2001. C195

[11] G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16:351–367, 2001. doi:10.1214/ss/1015346320 C196

[12] D. Rösch and H. Scheule. A multifactor approach for systematic default and recovery risk. *The Journal of Fixed Income*, 15(2):63–75, 2005. doi:10.3905/jfi.2005.591610 C200

[13] P. V. Shevchenko. Estimation of operational risk capital charge under parameter uncertainty. *The Journal of Operational Risk*, 3(1):51–63, 2008.
http://www.risk.net/journal-of-operational-risk/journal/
2160943/journal-operational-risk-volume-number-spring-2008
C196

[14] P. V. Shevchenko. *Modelling Operational Risk Using Bayesian Inference*. Springer, Berlin, 2011. C195

[15] P. V. Shevchenko and G. Temnov. Modeling operational risk data reported above a time-varying threshold. *The Journal of Operational Risk*, 4(2):19–42, 2009.
http://www.risk.net/journal-of-operational-risk/journal/
2160936/journal-operational-risk-volume-number-summer-2009
C195

Author addresses

1. **Pavel V. Shevchenko**, CSIRO Mathematics, Informatics and Statistics, Locked Bag 17, North Ryde, 1670, Australia
   mailto:pavel.shevchenko@csiro.au

2. **Xiaolin Luo**, CSIRO Mathematics, Informatics and Statistics, Locked Bag 17, North Ryde, 1670, Australia
   mailto:xiaolin.luo@csiro.au