Time matters:
How default resolution times impact final loss rates*

Jennifer Betz† Ralf Kellner‡ Daniel Rösch§

February 28, 2019

Abstract

Using access to a unique bank loss data base, we show positive dependencies of default resolution times (DRTs) of defaulted bank loan contracts and final loan loss rates (losses given default, LGDs). Due to this interconnection, LGD predictions made at the time of default and during resolution are subject to censoring. Pure (standard) LGD models are not able to capture effects of censoring and entail parameter distortions and, thus, an underestimation of average LGDs. In this paper, we develop a Bayesian hierarchical modeling framework for DRTs and LGDs which enables adequate unconditional LGD predictions and consistent LGD predictions conditional on the time in default in accordance with recent regulatory guidelines within one modeling framework. The proposed method is applicable to duration processes in general where the final outcomes depend on the duration of the process and are affected by censoring. By this means, potential parameter inconsistencies are overcome to ensure adequate predictions.

Keywords: default resolution time; loss given default; random effects

JEL classification: C23, G21, G33

* The authors would like to thank Global Credit Data (GCD) for granting access to their data base and the participants of NBER (time series) 2018 conference, the CREDIT 2018 conference, and the CEQURA 2018 conference for fruitful discussions and comments on an earlier version of this paper. The authors are particularly grateful to Prof. Alain Monfort for a thorough and valuable discussion of the paper during the CREDIT 2018 conference.

† Corresponding author, Universität Regensburg, Chair of Statistics and Risk Management, Universitätsstraße 31, 93040 Regensburg, Germany, email: jennifer.betz@ur.de

‡ Universität des Saarlandes, Chair of Quantitative Methods and Statistics, Campus C3 1, 66123 Saarbrücken, Germany, email: ralf.kellner@uni-saarland.de

§ Universität Regensburg, Chair of Statistics and Risk Management, Universitätsstraße 31, 93040 Regensburg, Germany, email: daniel.roesch@ur.de

1 Introduction

One of the most important tasks for financial institutions – such as banks and insurances companies – is the estimation and prediction of probabilities of default (PDs) and losses given default (LGDs) for loan contracts, whereby, the latter is a fraction of loss over the exposure at default. While PD estimation has a long lasting history, LGD estimation is in the focus of more recent research and regulation. In contrast to default itself, final losses of defaulted loan contracts are not observable until the end of the resolution process. These processes might continue for multiple years, thus, their duration (default resolution time, DRT) is subject to censoring. In an LGD modeling context, neglecting censoring in the resolution time entails parameter distortions and incorrect predictions. Due to positive dependencies of DRTs and LGDs, effects of censoring are shifted to the loss side (see Section 2). These effects are not captured in pure (standard) LGD models. The aim of this paper is to develop a joint modeling framework for DRTs and LGDs to analyse the dependence structure among DRTs and LGDs and diminish parameter distortions due to censoring and, thereby, enable adequate LGD predictions. Furthermore, this framework enriches LGD predictions for defaulted exposures by additional information in terms of the time that a loan contract has already spent in default (time in default) and, thus, allows for consistent predictions conditional on this time period.

Just recently, the European Banking Authority (EBA) published *Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures* (see European Banking Authority, 2017). The main aim of these guidelines is to reduce unjustified variability of risk parameters – such as PDs and LGDs. The EBA traces this variability back to deviations in the treatment of defaulted exposures among financial institutions. Besides the LGD for non-defaulted exposures (see Basel Committee on Banking Supervision, 2004), financial institutions are supposed to estimate LGD-in-default and an expected loss best estimate (EL_{BE}) for defaulted exposures. LGD-in-default should reflect economic downturn conditions in accordance with downturn LGDs for the non-defaulted exposures (see Basel Committee on Banking Supervision, 2004, 2005). In contrary, EL_{BE} is a point-in-time estimate reflecting the current economic surrounding. The guidelines demand consistent estimation methods regarding LGDs for non-defaulted and defaulted exposures (see European Banking Authority, 2017, §100) and the inclusion of post-default information – such as the time in default – for defaulted exposures (see European

Banking Authority, 2017, §168).

In recent years, the literature regarding LGD modeling has widened considerably. Comparative studies can be found in, e.g., Qi and Zhao (2011) and Loterman et al. (2012). However, literature considering *workout* LGDs is still limited. Most of the publications refer to *market-based* LGDs, whereby, the corresponding Recovery Rate (RR) is defined as ratio of the market price 90 days after default over the outstanding amount. Hence, market-based LGDs are only available for traded securities such as bonds. Workout LGDs are based on actual recovery payments collected during the resolution process and, thus, usually applied for loans. The distribution of workout LGDs is more extreme compared to market-based LGDs and, typically, high probability masses at no loss (LGD = 0) and total loss (LGD = 1) arise (see, e.g., Krüger and Rösch, 2017; Betz et al., 2018). Thus, the consideration of the distributional form is essential for workout LGDs. Altman and Kalotay (2014) develop a Bayesian Finite Mixture Model (FMM) with a probabilistic substructure in terms of an ordered logit (OL) model to estimate the probability of the mixture components depending on explanatory variables. A frequentistic version of this model is presented by Kalotay and Altman (2017). The model of Altman and Kalotay (2014) and Kalotay and Altman (2017) is applied by Bijak and Thomas (2015) and extended by Betz et al. (2018). Calabrese (2014) estimates a mixture of Beta distributions, whereas, Krüger and Rösch (2017) apply quantile regression on the LGD distribution. The literature regarding DRTs is more sparse and mainly refers to the duration of Chapter 7 and Chapter 11 resolutions (see, e.g., Helwege, 1999; Partington et al., 2001; Bris et al., 2006; Denis and Rodgers, 2007; Wong et al., 2007). Betz et al. (2016) and Betz et al. (2017) analyze DRTs of loan contracts and descriptively find impacts of DRTs on LGDs. The interconnection of DRTs and LGDs is also indicated in the related LGD literature. Dermine and Neto de Carvalho (2006) apply mortality analysis on a data set of bank loans, whereas, Gürtler and Hibbeln (2013) inter alia focus on the effects of censoring on LGD observations (see Section 2 for further information on effects of censoring). They do not provide any methodical solution, but suggest to restrict the data set to avoid biased estimates. However, LGD data is sparse so constraints might be unfavorable.

In this paper, we develop a hierarchical Bayesian modeling approach for joint estimation of DRTs and LGDs combining a finite mixture model (FMM) – which is well suited to capture the distributional features of workout LGDs (see Betz et al., 2018) – with a probabilistic substructure for the LGD and an accelerated failure time (AFT) model for the DRT. The inclusion of survival modeling techniques – in terms of the AFT model – in the LGD modeling context enables the consideration of censoring in LGDs. Thus, the hierarchical approach enables adequate unconditional LGD predictions for non-defaulted exposures and consistent LGD predictions conditional on the time in default for defaulted exposures within one modeling framework. Furthermore, correlated random effects are implemented in the hierarchical approach to allow for comovements of DRTs and LGDs in the time line. We apply the hierarchical approach to a unique European data set provided by Global Credit Data (GCD). GCD is a non profit initiative which aims to support banks to measure credit risk by collecting and analyzing historical loss data (see `http://www.globalcreditdata.org/` for further information). Furthermore, we compare the hierarchical approach to a pure (standard) LGD model in terms of an FMM with probabilistic substructure. By this means, we contribute to the literature in three ways. First, we deeply examine the dependence structure of DRTs and LGDs allowing for a direct and an indirect channel. We find positive impacts of DRTs on LGD distributions (direct channel) which are even more pronounced in boom and crisis periods (indirect channel). In crisis periods, this burdens financial market liquidity as more loan contracts are stuck in the resolution process. On top of that, losses are even higher than indicated by the direct channel due to an intensified dependence (indirect channel). Second, we illustrate consequences of neglecting censoring in an LGD modeling context. A pure (standard) LGD model suffers from parameter distortions which imply erroneous LGD predictions. Thus, LGDs are underestimated in an unconditional perspective by up to 20 percentage points. Considering defaulted exposures, i.e., LGD-in-default and EL_{BE}, this underestimation is intensified. Third, the hierarchical approach diminishes parameter distortions and, thus, leads to adequate unconditional LGD predictions. Furthermore, it offers an intuitive framework to generate LGD predictions conditional on the time in default, i.e., LGD-in-default and EL_{BE}, in accordance with the EBA guidelines (see European Banking Authority, 2017). At first glance, the described application seems to be rather specific in credit risk management. However, the proposed modeling framework is applicable to duration processes in general where dependence between time and some outcome at the end of the time horizon is present.

The remainder of this paper is structured as follows. Section 2 provides further information on the effects of censoring on LGDs and, thus, reasoning for the introduction of the EBA guidelines. Section 3 introduces the hierarchical modeling approach. Data and results are presented in

Section 4. In Section 5, the model is validated on an in sample and out of sample perspective. Section 6 concludes.

2 Background

As stated in Section 1, *workout* LGDs are characterized by time specific censoring, as they are indirectly affected by the DRT. This section gives an impression of the effects of censoring on LGDs and, thus, offers reasoning for the EBA guidelines with respect to the treatment of defaulted exposures (see European Banking Authority, 2017).

The default resolution process of corporate loan contracts is rather complex as it depends on local insolvency codes, common business practices, and individual decisions. The parameters of the process, i.e., the DRT and the LGD, are affected by this complexity. Simpler – or more efficient – default resolutions might be described by short DRTs and low LGDs, while higher complexity results in longer resolutions with higher losses. Figure 1 illustrates the scheme of resolution processes. The process starts with the default of a debtor, i.e., its loan contract, and

Figure 1: Structure of default resolution

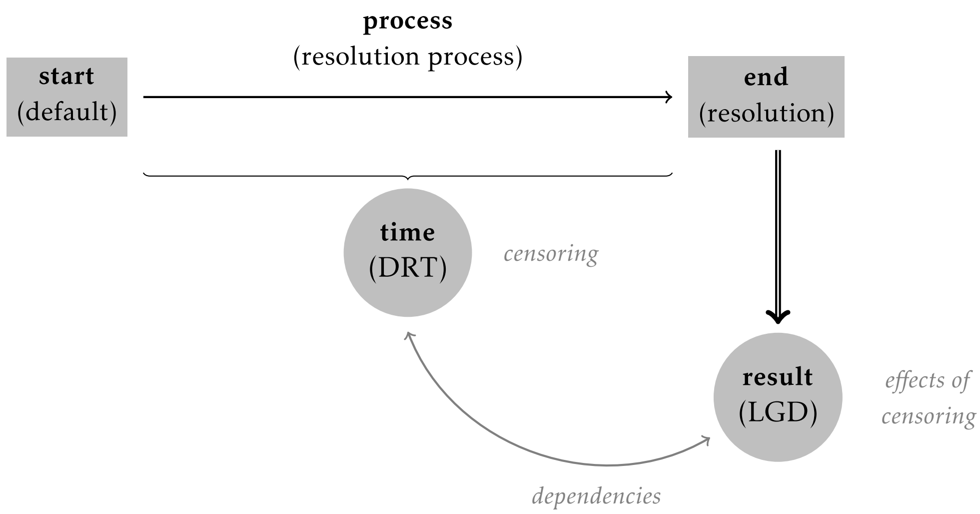

Note: The figure illustrates a schematic representation of the resolution process. The process is limited by its start point (default) and its end point (default resolution) presented by gray rectangles. It can be characterized by its parameters time (DRT) and result (LGD) presented by gray circles. The time is subject to censoring. The result variable is measurable on metric scale. The gray double arrow between the parameters indicate dependence.

ends with default resolution. The duration of this process differs among contracts and is subject to censoring. In particular, longer DRTs are not observable until the end of the observation period and are censored to the time in default. Due to the dependency between DRTs and LGDs, effects of censoring can be remarkable on the LGD side. As the final DRTs of censored observations are longer on average, final LGDs tend to be higher (but not observable). This phenomenon is often refereed to as resolution bias.

Generally, neglecting censoring might result in biased parameter estimates. We illustrate this by a simplified example in order to ease the understanding of censoring effects. For instance, assume a Weibull distribution for DRTs ($T \sim$ Weibull) and – for simplicity – a fixed censoring time c. To ease transparency, we further assume that LGDs linearly depend on DRTs, thus, following the linear model

$$L_i = \alpha + \beta \ \ln(T_i) + e_i \ , \tag{1}$$

where L_i is the LGD of loan i and $e_i \sim N(0, \sigma^2)$. We randomly generate 10,000 pairs of DRTs and LGDs by Monte-Carlo simulation based on this model. We choose the true values of the parameters to match location and scale of empirical DRTs and LGDs. The shape and scale of the Weibull distribution are set to one and two to ensure an average DRT of round about two years. In Equation (1), we set $\alpha = 0.2$ and $\beta = 0.3$ resulting in an average LGD of round about 0.25. Please note that the choice of parameters is arbitrary as resulting effects are independent of the specific values. Figure 2 illustrates the impact of censoring on the result variable LGD. The light gray bars illustrate the simulation result for the *true* distributions (which are not observable), whereas, the dark gray bars are the observable censored distributions. The lines display the corresponding mean values, whereby, the dotted lines indicate absence of censoring. In the upper panel, the distribution of the time is displayed. Censoring limits the distribution to a certain value which implies a slight underestimation of the average value if censoring is neglected. Due to the linear dependence between time and result as of Equation (1), this phenomenon is transferred to the result variable (see lower left panel). Thus, average values are underestimated assuming a positive dependence of time and result as the true mean marked by a dotted line is above the mean in the presence of censoring. In the conditional perspective, only censored cases are considered. Hereby, the effect of censoring and, thus, the underestimation, is intensified (see lower right panel) as the difference between true mean and mean in the presence of censoring increases. Neglecting censoring implies the same distribution for censored values as for non-censored ones. However, the true distribution is shifted towards higher values assuming a positive dependence, i.e., $\beta > 0$.

Figure 2: Impact of censoring on the result variable

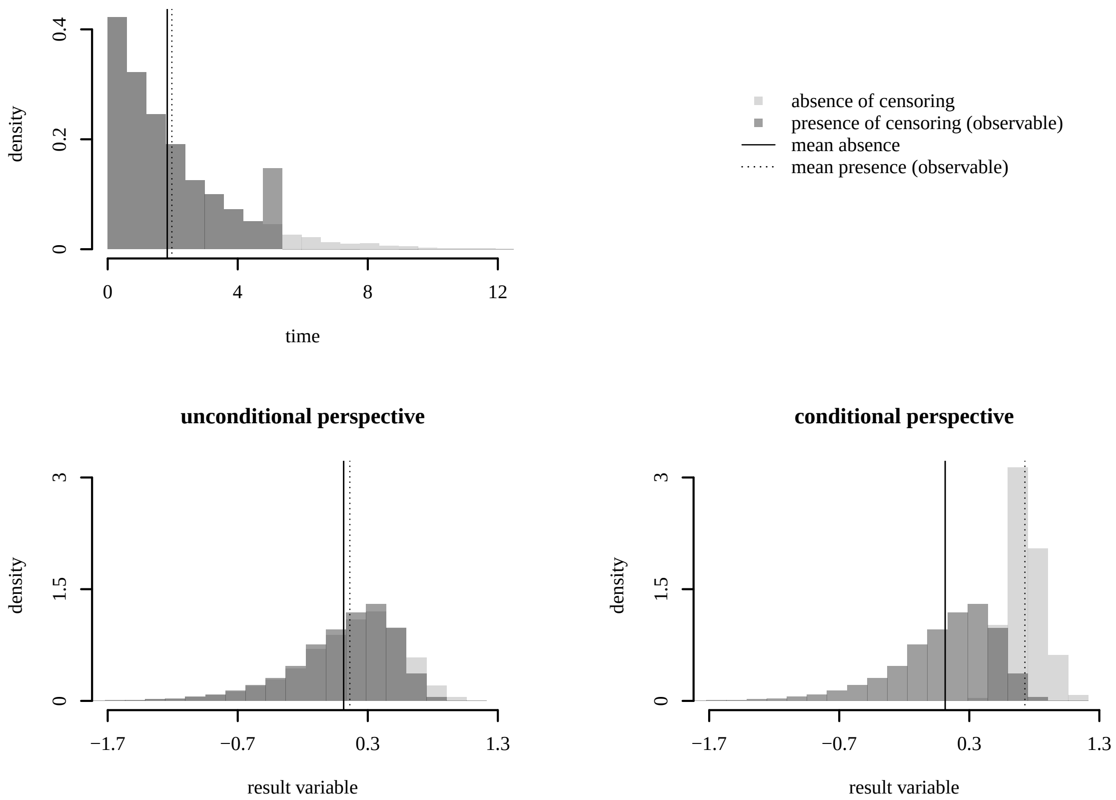

Note: The figure illustrates the impact of censoring in the time variable on the result variable. The upper panel displays the distribution of time in the absence of censoring (light gray bars) and in the presence of censoring (dark gray bars). The distribution in the absence of censoring might be interpreted as the *true* distribution. The lower left panel illustrates the unconditional distribution of the result variable (dark gray indicates censoring, light gray no censoring). The lower right panel restricts the presentation to the censored cases. Means of the distributions are marked by lines, whereby, the dotted line indicates absence of censoring.

In the context of the EBA guidelines, the unconditional perspective corresponds to unconditional LGDs for the non-defaulted exposure, whereas, the conditional perspective reflects LGDs conditional on the time in default for the defaulted exposure. Thus, the consideration of post-default information – such as the time in default – is required. The hierarchical approach we develop is a joint modeling approach for DRTs and LGDs. It considers censoring, dependencies between DRTs and LGDs and, thus, allows for adequate unconditional and consistent conditional LGD predictions within a single modeling framework.

At the first glance, the described setting seems to be a rather special case in credit risk management. However, applications are diverse. Generally, duration processes are subject to censoring. Whenever time dependent result variables on a metric scale are of main interest, censoring should be considered to avoid underestimation (or overestimation respectively, if negative dependencies between time and result are present, i.e., $\beta < 0$). Examples might be found in business where complex negotiations lead to outcomes on a metric scale, e.g., granting loans (negotiation process vs. granted amount). Besides, applications in completely different fields are conceivable, e.g., medicine and health science (healing process vs. resulting quality or strength). The developed hierarchical approach might be adjusted to different applications regarding the characteristics of the outcome variable. In this paper, we apply an FMM to model the distribution of LGDs which is rather complex. In a setting where the result variable follows a less extreme distribution, the FMM can be exchanged by a much simpler model.

3 Methods

This paper develops a hierarchical modeling approach for DRTs and LGDs and, thereby, analyzes the dependence structure of DRTs and LGDs and reveals impacts of censoring in DRTs on LGDs. We combine a pure (standard) LGD model with survival analysis techniques – in terms of an accelerated failure time (AFT) model – in a hierarchical structure to consider censoring in an LGD modeling context. As stated in Section 1, workout LGDs are characterized by a rather extreme distributional form. The distribution is bimodal with high probability masses at no and total loss. Furthermore, the two modes are characterized by bindings, i.e., values which are exactly 0 or exactly 1. We, therefore, extend the Bayesian finite mixture model (FMM) with a probabilistic substructure in terms of an ordered logit (OL) model developed by Altman and Kalotay (2014) and further extended by Betz et al. (2018). This model seems to be well suited to capture the characteristic features of LGD distributions. To investigate direct dependencies of DRTs on LGDs, the DRT serves as an explanatory variable in the LGD model. Two correlated random effects are included to study comovements of DRTs and LGDs in the time line (indirect dependencies). In the following, we briefly review the LGD model of Altman and Kalotay (2014) and Betz et al. (2018) and discuss the extensions in the context of the hierarchical model.

LGD Model

We apply a Normal FMM to model the distribution of LGDs. Generally, FMMs offer high flexibility in modeling distributions of unknown shape (see McLachlan and Peel, 2000). The dependent variable L is assumed to be divided into a finite number of K latent classes. In each class k, L follows a Normal distribution with parameters θ_k depending on the latent class k. We assume normally distributed components to achieve computational transparency (see McLachlan and Peel, 2000). Thus, the probability density function (PDF) of an FMM $g(L|\theta_1,\dots,\theta_K)$ is the p_k weighted sum of the component PDFs $f_k(L|\theta_k)$:

$$g(L|\theta_1,\dots,\theta_K) = \sum_{k=1}^{K} p_k \; f_k(L|\theta_k)\,. \tag{2}$$

To ensure the general properties of a PDF, i.e., $g(l) \geq 0$ for all $l \in \mathbb{R}$ and $\int_{-\infty}^{\infty} g(l) = 1$, $p_k \geq 0$ and $\sum_k p_k = 1$ must hold. Assuming conditional independence, the likelihood of a Normal FMM $\phi(L_1, ..., L_N \,|\, \mu, \sigma, p)$ is the product of the individual likelihood contributions:

$$\phi(L_1,...,L_N \,|\, \mu,\sigma,p) = \frac{1}{(2\pi)^{\frac{N}{2}}} \prod_{i=1}^{N} \left(\sum_{k=1}^{K} \frac{p_k}{\sigma_k} \exp\left[-\frac{(L_i-\mu_k)^2}{2\sigma_k^2} \right] \right), \tag{3}$$

where, μ_k and σ_k are the parameters of the latent class k and N is the number of observations. To adapt data augmentation, the component weight p_k is replaced by an indicator variable d_{ik} which takes the value one if l_i is a random draw of component k and zero otherwise:

$$\phi(L_1,...,L_N \,|\, \mu,\sigma,d) = \frac{1}{(2\pi)^{\frac{N}{2}}} \prod_{i=1}^{N} \left(\sum_{k=1}^{K} \frac{d_{ik}}{\sigma_k} \exp\left[-\frac{(L_i-\mu_k)^2}{2\sigma_k^2} \right] \right). \tag{4}$$

To identify loan contracts with no and total loss, we fix the parameters of the two outer components. The means are set to $\mu_1 = 0$ and $\mu_K = 1$ with small standard deviations ($\sigma_1 = \sigma_K = 0.0001$). Results are robust to different (reasonable) values for σ_1 and σ_K.

To estimate the probability for loan i of belonging to the k-th component depending on covariates, a probabilistic substructure in terms of an OL model is formulated. To rely on the classical formulation of the OL model, we define the component affiliation y_i:

$$y_i = k \quad \text{if } d_{ik} = 1\,, \tag{5}$$

where, d_{ik} is the indicator as of Equation (4). The component affiliation Y_i is categorically distributed and determined by the location of a metric latent variable Y_i^* to corresponding cut points c_k ($k \in \{1,\ldots,K-1\}$):

$$Y_i = \begin{cases} 1 & \text{if } Y_i^* \leq c_1 \\ 2 & \text{if } c_1 < Y_i^* \leq c_2 \\ \vdots & \\ K & \text{if } c_{K-1} < Y_i^* \,. \end{cases} \tag{6}$$

The latent variable Y_i^* follows a linear model:

$$Y_i^* = z_i\zeta + F_{t(i)} + e_i, \qquad e_i \sim \text{logistic}\,, \tag{7}$$

where, z_i is a $(1 \times J)$ vector of independent variables and ζ is the $(J \times 1)$ vector of coefficients. The term e_i describes the errors. A random effect $F_{t(i)}$ is introduced into the modeling framework to control for comovement in the time line. It originates from a Normal distribution with mean zero and standard deviation σ:

$$F_t \sim \mathrm{N}(0, \sigma)\,. \tag{8}$$

The time stamp $t(i)$ in Equation (7) indicates the default time t in quarters of loan i. Two loans i and i' which defaulted in the same quarter ($t(i) = t(i') = t$) share the same realization of the random effect ($f_{t(i)} = f_{t(i')} = f_t$). For $f_t > 0$ ($f_t < 0$), both loans exhibit higher (lower) values of y_i^* and, thus, higher (lower) probabilities of high component affiliations y_i. Higher component affiliations y_i are accompanied with higher loss rates and vice versa. Thus, the random effect displays the comovement in time line, i.e., higher or lower average loss rates in specific default quarters which can not be explained by observable variables included in z_i.

Betz et al. (2018) additionally consider an autoregressive process of order 1, i.e., AR(1), for the random effect to allow for cyclical movements in the realizations of the random effect. In this paper, we do not consider this specification due to simplicity as the specification of the random effect seems to have negligible impact on its realizations. For conditional predictions (LGDs-in-default, $\mathrm{EL}_{\mathrm{BE}}$) in the hierarchical approach we apply the realized value of the random effect in the corresponding time period, while we use the mean value 0 for unconditional predictions.

Hierarchical Model

In the hierarchical model, the pure (standard) LGD model of the previous paragraphs is extended by an additional hierarchical level in terms of an AFT model for the DRT to consider censoring and allow for LGD predictions conditional on the time in default. The logarithm of the resolution time $\ln(T_i)$ can be expressed by a linear model:

$$\ln(T_i) = \beta_0 + x_i\beta + F^T_{t(i)} + s\,\varepsilon_i\,, \qquad \varepsilon_i \sim \text{negative Gumbel}\,, \tag{9}$$

where x_i is a $(1 \times J_T)$ vector of independent variables, β is the $(J_T \times 1)$ vector of coefficients, and β_0 is the intercept. We assume the errors ε_i to follow a negative Gumbel distribution and, thus, the DRT to be Weibull distributed. Different distributional assumptions for the errors are possible with the Normal, Logistic, Exponential and Weibull distribution being the most common ones in the AFT model. Among these, the Weibull distribution seems to have the best fit. The term s is the scale parameter. A random effect $F^T_{t(i)}$ is introduced into the modeling framework to control for comovement in the time line. Equation (9) applies to non censored, i.e., final, observations. For censored observations, final realizations are estimated within the Bayesian modeling framework. By this means, we are able to predict final DRTs for censored data points, i.e., unresolved loans.

In the hierarchical approach, the AFT model for the DRT is simultaneously estimated with FMM for the LGD (see previous paragraph). To develop an intuitive method to generate LGD predictions conditional on the time in default, the logarithm of the DRT is included as explanatory variable. This accounts for direct dependencies of DRTs and LGDs. Equation (7) modifies to:

$$\mathcal{Y}^*_i = z_i\gamma + \ln(T_i)\gamma_T + F^L_{t(i)} + \epsilon_i\,, \qquad \epsilon_i \sim \text{logistic}\,, \tag{10}$$

where z_i is the $(1 \times J_L)$ vector of independent variables, γ is the $(J_L \times 1)$ vector of coefficients, and γ_T is the coefficient of the logarithm of the DRT. Again, a random effect $F^L_{t(i)}$ is introduced into the modeling framework to control for comovement in the time line. Equation (2), (3), (4), (5), and (6) apply in analogy to Y_i*.

The random effects $F^T_{t(i)}$ as of Equation (9) and $F^L_{t(i)}$ as of Equation (10) originate from a bivariate Normal distribution:

$$\begin{pmatrix} F^T_t \\ F^L_t \end{pmatrix} \sim \mathrm{N}_2\left(0_2\,, \Sigma\right)\,, \tag{11}$$

where 0_2 is the two dimensional zero vector ($0_k = (\begin{array}{cc} 0 & 0 \end{array})^T$) and Σ is the (2×2) covariance matrix. The latter is based on the individual standard deviations (σ_T and σ_L) and the (2×2) correlation matrix Ω:

$$\begin{aligned} \Sigma &= \operatorname{diag}(\sigma_T, \sigma_L)\ \Omega\ \operatorname{diag}(\sigma_T, \sigma_L) \\ &= \begin{pmatrix} \sigma_T^2 & \sigma_T \sigma_L \omega_{L,T} \\ \sigma_T \sigma_L \omega_{T,L} & \sigma_L^2 \end{pmatrix}, \end{aligned} \tag{12}$$

where $\omega_{T,L} (= \omega_{L,T})$ is the correlation of F_t^T and F_t^L. By the inclusion of the random effects, we control for joint comovements of loss rates and resolution times in the time line. Two loans i and i' which defaulted in the same quarter ($t(i) = t(i') = t$) share the same realizations of the random effects ($f_{t(i)}^T = f_{t(i')}^T = f_t^T$ and $f_{t(i)}^L = f_{t(i')}^L = f_t^L$, however, $f_t^T \neq f_t^L$ in most of the cases). For $f_t^T > 0$ ($f_t^T < 0$), average DRTs are higher (lower). Assuming a positive correlation between the random effects and a positive parameter estimate of the logarithm of the DRT in the LGD model ($\gamma_T > 0$), the corresponding LGDs are effected through two channels: Directly, as higher (lower) DRTs are inserted in the LGD model. Indirectly, as positive (negative) realizations of f_t^T tend to imply positive (negative) realizations of f_t^L due to the positive correlation. Thus, LGDs are even higher. However, negative realizations of f_t^L remain possible for a stochastic process as of Equation (11) which might reduce LGDs. Both scenarios are conceivable. Confronted with a tense economic surrounding, financial institutions might decide to follow a wait-and-see strategy and relocate resolution efforts in the future. This might provide benefits and reduce the LGD ($f_t^L < 0$). However, LGDs might be further increased ($f_t^L > 0$) if financial institutions are forced to resolve defaulted loans at a certain point in time, e.g., if there is no further option to wait.

Estimation

The parameters of the LGD model and the hierarchical model are estimated via Bayesian inference. Following Bayes' theorem, the posterior distribution of the parameters θ given the data is

$$f(\theta | \text{data}) = \frac{f(\text{data} | \theta)\, f(\theta)}{f(\text{data})}, \tag{13}$$

where $f(\theta|\text{data})$ is the posterior distribution of the parameters, $f(\text{data}|\theta)$ the likelihood of the data given the parameters θ, and $f(\theta)$ the prior distribution of the parameters. The denominator $f(\text{data})$ guarantees the properties of a density function and is calculated as $f(\text{data}) = \int f(\text{data}|\theta)\, f(\theta)\, d\theta$. In many cases, the integral in the denominator can not be solved analytically. Markov Chain Monte Carlo (MCMC) methods might be used to sample from the posterior distribution. Most commonly, Metropolis Hastings (MH) and Gibbs (as special case of MH) algorithm are applied. Confronted with hierarchical models and, thus, highly correlated posterior distributions, both algorithms suffer from an inappropriate high number of iterations to reach the equilibrium distribution. Therefore, we apply the Hamiltonian Monte Carlo (HMC) algorithm. The HMC algorithm is efficiently implemented in Stan (see Stan Development Team, 2016, for further information on the implementation) and combines a Hamiltonian evolution with a Metropolis proposal to reduce the correlation in the chains. The parameters of the proposal distribution and the Hamiltonian evolution are tuned during the adaption phase. The LGD model and the hierarchical model are sampled with two HMC chains. Burn-in is set to 500. Posterior samples contain 25,000 iterations per chain with a thinning of 5. Metric dependent variables are standardized to ease convergence. Most of the model parameters are provided with weakly informative prior distributions. See Section 1 of the online companion for this paper for detailed information on the Bayesian model specifications. Common convergence diagnostics can be found in Section 2 of the online companion.

The uncertainty in the parameter estimates expressed by the dispersion in the posterior distributions $f(\theta|\text{data})$ should be reflected in predictive distributions. Given a new data point $\tilde{a}$ the posterior predicted distribution is

$$f(\tilde{a}|\text{data}) = \int_{\Theta} f(\tilde{a}|\theta, \text{data})\, f(\theta|\text{data})\, d\theta\,, \tag{14}$$

where $f(\tilde{a}|\theta, \text{data})$ is the predictive distribution for $\tilde{a}$ given one single parameter estimate θ and $f(\theta|\text{data})$ is the posterior distribution of that parameter. The quantity Θ (with $\theta \in \Theta$) denotes the parameter space. Inserting each iteration of θ from $f(\theta|\text{data})$ based on the HMC algorithm into the modeling framework corresponds to sampling from $f(\tilde{a}|\text{data})$.

4 Data and results

4.1 Data

We use access to the unique loss data base of Global Credit Data (GCD). The data base includes detailed loss information on transaction basis of 53 member banks all around the world. In the data base, the LGD is determined by $L_i = 1 - \text{RR}_i$, whereby, L_i is the loss rate of loan i and RR_i is the corresponding recovery rate (RR). The RR is calculated as the sum over the present values of all relevant transactions divided by the outstanding amount (see Betz et al., 2016, 2018).

We follow Höcht and Zagst (2010) and Höcht et al. (2011) and develop two selection criteria to eliminate loans with extraordinary payment structures. Both criteria relate all relevant transactions including charge-offs (which are not included in the LGD calculation) to the outstanding amount. The first criterion, to which we refer as *pre-resolution criterion*, relates transactions arising pre-resolution to the outstanding amount at default. We set the barriers of the pre-resolution criterion to $[90\%, 110\%]$ for resolved and $[-50\%, 400\%]$ for unresolved loans. In the second criterion, i.e., *post-resolution criterion*, transactions occurring post resolution are related to a fictional outstanding amount at resolution. The barriers are set to $[-10\%, 110\%]$ for the post-resolution criterion. The post-resolution criterion applies for resolved loans only. Subsequently, loans with abnormal low and high LGDs ($< -25\%$ and $> 125\%$) are eliminated. Overall, 0.50% of resolved loans are eliminated due to the pre-resolution criterion and 0.19% due to the post-resolution criterion, whereas, 0.23% of unresolved loans are eliminated based on the pre-resolution criterion. Subsequently, 0.13% are sorted out due to abnormal low and high LGD values. We consider a subsample of defaulted European term loans and lines to small and medium sized enterprises (SMEs). The twelve most common European countries in the data base – i.e., Great Britain, Germany, Denmark, Portugal, Ireland, France, Finland, Sweden, Norway, Latvia, Estonia, and Poland – are included. We further exclude loans which defaulted before 2004 and after 2016 (10.02% of subsample data). A subsample of 38,165 loans remains.

Figure 3 illustrates the interconnection of the two parameters of the resolution process, i.e., the DRT and the LGD. Therefore, the data set is divided into DRT buckets based on realized DRTs. The first bucket includes all loans with DRTs $\in [0, 0.5]$ years. The second bucket contains all loans with DRTs ∈ (0.5, 1.0] years, and so on (see x-axis of left panel and legend of right panel). In the left panel, box plots of LGDs divided by DRT buckets are displayed. The thick black lines

Figure 3: Relation of DRT and LGD

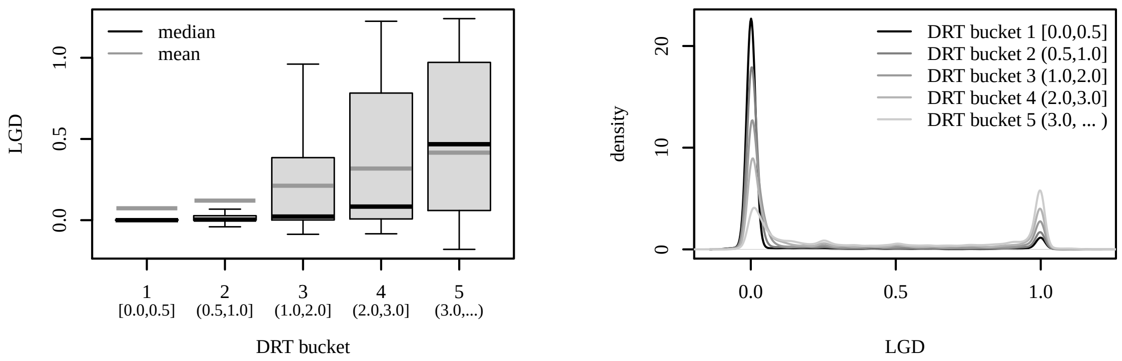

Note: The figure illustrates the relation of DRTs and LGDs. The data is divided into DRT buckets based on the realized DRTs. Thus, the first bucket includes all loans with DRTs ∈ [0, 0.5] years. The second bucket contains all loans with DRTs ∈ (0.5, 1.0] years, and so on (see x-axis of left panel and legend of right panel). In the left panel, box plots of LGDs for the DRT buckets are displayed. Outliers are hidden. The thick black lines mark the medians, whereas, the thick gray lines are the means. In the right panel, kernel density estimates of LGDs for the DRT buckets are illustrated. The band width is fixed to 0.015 to ensure comparability.

mark the medians, whereas, the thick gray lines are the means. Considering the latter, average LGDs seem to linearly increase in the DRT buckets. To examine the origin of this increase, the right panel displays kernel density estimates for the DRT buckets. The LGD distribution of higher DRT buckets is shifted towards higher LGD values, i.e., probability masses of lower LGD values decrease and probability masses of higher LGD values increase. Thus, average values increase.

Table 1 summarizes the descriptive statistics of the dependent and explanatory variables. Figures are stated for all loans (resolved, i.e., non censored, and unresolved, i.e., censored, cases) and for resolved and unresolved loans separately. The upper panel of the table includes descriptive statistics for the LGD and the DRT. For unresolved cases, incurred LGDs are applied. Incurred LGDs are computed as the sum over the present values of all relevant transactions, which occurred up to the end of the observation period (end of 2016), divided by the outstanding amount. As the resolution process is not terminated, incurred LGDs are higher than final LGDs. DRTs for unresolved cases are censored to the end of the observation period (end of 2016), e.g., for unresolved loans defaulted at the end of 2015, a censored DRT of one year is assigned. Censored DRTs are lower than final DRTs as the resolution process

Table 1: Descriptive statistics

		all	**resolved**	**unresolved**
n		38,165	35,272	2,893
		dependent variables		
LGD	mean	0.2534	0.2099	0.7839
	median	0.0133	0.0082	0.9780
	standard deviation	0.3810	0.3531	0.3017
DRT	mean	1.9882	1.7342	5.0839
	median	1.2621	1.1335	4.9090
	standard deviation	2.0756	1.7509	3.0147
		loan specific (metric)		
EAD	mean	533,118.89	516,582.05	734,739.20
	median	102,987.29	100,237.48	155,097.53
	standard deviation	3,624,711.52	3,610,978.60	3,782,983.43
		loan specific (categoric)		
Facility	*term loan*	62.00%	60.39%	81.68%
	line	38.00%	39.61%	18.32%
Protection	*no*	25.61%	26.03%	20.46%
	yes	74.39%	73.97%	79.54%
Industry	*non FIRE*	83.13%	82.13%	95.30%
	FIRE	16.87%	17.87%	4.70%
		macro variables		
Δ HPI	mean	-1.6966	-1.7765	-0.7229
	median	0.1662	0.1662	0.8360
	standard deviation	6.0221	5.9928	6.2878
VIX	mean	24.4762	24.3681	25.7947
	median	22.9249	22.6771	23.3451
	standard deviation	9.4105	9.4699	8.5465

Note: The table summarizes descriptive statistics for dependent and independent variables in the data set. For metric variables, means, medians, and standard deviations are stated. Proportions are presented for variables of categoric nature. The sample size is denoted by *n*. The abbreviation FIRE means *Finance*, *Insurance*, *Real Estate* and denotes corporations of this industries. The macro variable Δ HPI is the is the yoy percentage chance of the *House Price Index*, whereas, the VIX is the *Volatility Index*.

is not terminated. In the table, average values of LGDs and DRTs for unresolved cases are higher compared to resolved cases as unresolved cases are shaped by rather *bad* loans, i.e., loans exhibiting high DRTs and high LGDs. In the middle panels of the table, descriptive statistics of loan specific independent variables are stated. We use the EAD to control for the size of the loan. It is further distinguished between term loans and lines, whether a loan is protected by collateral or guarantee or not, and whether the debtor has Finance, Insurance, Real Estate (FIRE) industry affiliation. Reference categories in the subsequent models are printed italic in the table. The lower panel of the table contains descriptive statistics of the applied macroeconomic variables. The year-on-year (yoy) percentage change of weighted average real residential prices (Δ HPI) is employed as explanatory variable for the LGD, whereas, we use the VSTOXX Volatility Index (VIX) for the DRT. We tested further macro variables, e.g, the yoy percentage change of weighted average seasonally adjusted GDPs and the quarterly average yoy percentage change of weighted average equity indices. However, Δ HPI and VIX exhibit the highest statistical evidence.

Figure 4 illustrates the time patterns of average DRTs in the left panel and average LGDs in the right panel for resolved loans (thick black line) and all loans (resolved and unresolved loans, thin gray line). Regarding the latter, values for unresolved loans, i.e., censored observations, have to be calculated. Thus, DRTs are censored to the end of the observation period (end of 2016) and incurrent LGDs are considered for unresolved cases. Incurrent LGDs are computed as sum over the present values of all relevant transactions, which occurred up to the end of the observation period (end of 2016), divided by the outstanding amount. The relation of DRTs and LGDs (see Figure 3) might partly be driven by analogous time patterns. Both dependent variables sharply increase prior to the Global Financial Crisis (GFC, 2007 Q2) and reach the maximum during the climax of the GFC. The rebound in the aftermath of the crisis seems gradual. There are only minor deviations between resolved loans and all loans considering the average DRTs. The graph for all loans is slightly shifted upwards by the censored observations. Regarding average LGDs, this spread is severe particularly in the most recent time periods. This is mainly due to incurrent LGDs, i.e., LGDs based on transactions which occur up to the end of the observation period, in the averaging. Final LGDs will be lower. However, final LGDs of all loans will still lie above the black line (final LGDs of resolved loans). Due to the effects of censoring, final LGDs are only observable for defaults with short DRTs in the more recent time

Figure 4: Time patterns of average DRTs and average LGDs

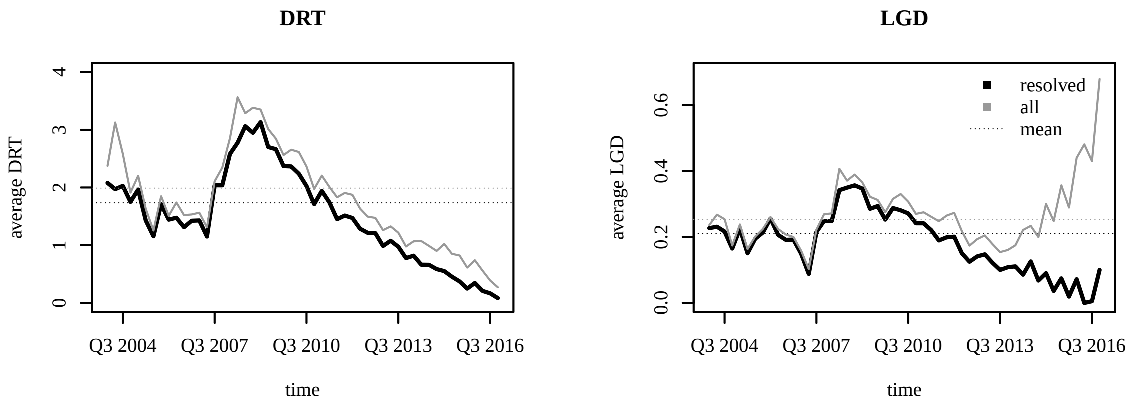

Note: The figure illustrates time patterns of average DRTs in the left panel and average LGDs in the right panel. The black lines display the average values for resolved loans, whereas, the gray lines are average values for all loans, i.e., resolved and unresolved cases. Thus, the latter include censored values. Means over the entire time period are illustrated by dotted lines.

periods (see Section 2). Due to the interconnection of DRTs and LGDs, these tend to be lower implying an underestimation of LGDs in the more recent time periods.

In this paper, we analyze the effects of censoring on unconditional and conditional LGD predictions on an in-sample and out-of-sample perspective. Therefore, we divide the data set as of Table 1 into subsamples. The first subsample serves as *estimation sample*. It includes all loans defaulted between 2004 Q1 and 2010 Q4. Thus, it comprises times of rather sound economic surrounding, the GFC, and parts of the rebound phase. We treat loans which are not resolved until 2010 Q4 as censored observations, i.e., unresolved loans. The second subsample, to which we refer to as *validation sample I*, includes the final observations to the censored observations as of the estimation sample. We apply validation sample I to perform an *out-of-sample validation* of LGDs. The third sample, i.e., *validation sample II*, includes all loans defaulted between 2011 Q1 and 2016 Q4. It is used to perform an *out-of-sample out-of-time validation* of LGDs. Table 2 summarizes the estimation sample and the validation samples. In the upper panel, the sample sizes are stated. Validation sample I consists of the 10,171 loans which are treated as unresolved cases in the estimation sample, i.e, are unresolved until the end of 2010. Some of these loans (1,724) are still unresolved at the end of 2016. However, the proportion of unresolved loans is lower in validation sample I compared to the estimation sample. In the lower panel, average values of LGDs and DRTs are stated. These are rather similar comparing the estimation sample and validation sample II, however, considerably higher in validation

Table 2: Estimation sample and validation samples

		estimation sample	**validation sample I** (*out-of-sample*)	**validation sample II** (*out-of-sample out-of-time*)
n	all	31,988	10,171	6,177
	resolved	21,817	8,447	5,008
	unresolved	10,171	1,724	1,169
		dependent variables		
average LGD	all	0.2586	0.4270	0.2267
	resolved	0.1801	0.3511	0.1017
	unresolved	0.4270	0.7987	0.7622
average DRT	all	1.5763	4.2566	1.0495
	resolved	1.1964	3.6851	0.7869
	unresolved	2.3911	7.0568	2.1743

Note: The table summarizes the applied samples. The number n, the average LGD, and the average DRT of all loans, resolved loans, and unresolved loans are presented for the estimation sample and the two validation samples. The models are estimated based on the estimation sample. This sample includes all loans defaulted between 2004 Q1 and 2010 Q4. Loans which are not resolved until 2010 Q4 are treated as censored observations, i.e., unresolved cases, in the estimation. Validation sample I contains the final observations of these unresolved cases. However, observations exist which are still censored at the end of 2016 (unresolved cases in validation sample I). In validation sample II, loans which defaulted between 2011 Q1 and 2016 Q4 are included. Thus, validation sample I is applied for the *out-of-sample* validation, whereas, the *out-of-sample out-of-time* validation is performed on validation sample II.

sample I. This is due to the fact that validation sample I contains final observations to censored cases in the estimation sample, thus, observations with higher DRTs and higher LGDs.

4.2 Results

In Bayesian inference, posterior distributions of parameters are assumed to be continuous. Thus, a single value of the posterior distribution has a probability of zero. On the contrary, one *single* parameter estimate is assigned in frequentistic terms. A null hypothesis is set up to reach a yes-or-no decision, e.g., whether the parameter is positive or not. Under the Bayesian approach, estimates are provided by posterior distributions which offer an intuitive consideration of parameter uncertainty (uncertainty refers to the range of possible parameter values and their probabilistic substructure, not to Knightian uncertainty according to which no probabilities can be assigned to parameter values). Thus, other concepts are indicated to quantify if results are in favor of an impact, i.e., if an impact is statistical evident. We apply two of them – *credible intervals* and *Bayes factors*.

Credible intervals, e.g., Highest Probability Density Intervals (HPDIs), specify intervals in the domain of the posterior distribution in which the unobservable parameter lies with a certain probability. The HPDI denotes the narrowest credible interval. If $0 \notin$ HPDI, the domain of the posterior distribution is located in the positive (negative) value range indicating statistical evidence of the positive (negative) sign. Besides credible intervals, we apply Bayes factors to evaluate statistical evidence. Bayes factors are the relation of posterior odds to prior odds. Posterior odds are defined as the ratio of posterior probability masses under the null hypotheses and the alternative hypothesis. As we are interested in the evidence of the signs, posterior odds are derived as the ratio of posterior mass favoring the sign of the posterior mean to posterior mass of the opposite sign:

$$\text{posterior odds}_{\mathrm{E}[\theta]<0} = \frac{\mathbb{P}(\theta < 0 \mid \text{data})}{\mathbb{P}(\theta \geq 0 \mid \text{data})}$$
$$\text{posterior odds}_{\mathrm{E}[\theta]>0} = \frac{\mathbb{P}(\theta > 0 \mid \text{data})}{\mathbb{P}(\theta \leq 0 \mid \text{data})},$$

whereby, θ denotes an arbitrary parameter. Assuming a positive posterior mean ($\mathrm{E}[\theta] > 0$), posterior $\text{odds}_{\mathrm{E}[\theta]>0} = 3$ indicates that a positive impact is three times as likely as a negative impact. Prior odds are the corresponding ratio of the prior distribution. Assuming a symmetric prior distribution around zero, posterior odds are equivalent to the Bayes factor. A Bayes factor exceeding 3.2 is deemed as substantial evidence. Values above 10 are assigned with strong evidence, whereas, values above 100 are related to decisive evidence (see Kass and Raftery, 1995).

LGD Model

The LGD model is estimated based on the estimation sample (see Table 2). However, it offers no possibility to include censored observations, i.e., unresolved loans, in the estimation process. Thus, the 21,817 resolved cases are included, whereas, 10,171 unresolved defaults are neglected. As these unresolved loans tend to exhibit higher LGDs due to the resolution bias, parameter estimates might be distorted.

Table 3 summarizes the results of the LGD model. Parameters are stated in the first column, whereas, the second column presents posterior means. In the FMM within the LGD model, parameters of the outer components (μ_1 and σ_1 for the first component, μ_5 and σ_5 for the fifth component) are fixed to identify loans with no (LGD = 0) and total (LGD = 1) loss. The second and third component are located nearby the first component ($\mu_2 = 0.0067$ and $\mu_3 = 0.0290$)

Table 3: Results of the LGD model

	posterior mean	HPDI (95%)		posterior odds	naive standard error	time series standard error
			LGD model			
μ_1	0.0000			*not estimated*		
μ_2	0.0067	0.0064	0.0070	∞	0.0000	0.0000
μ_3	0.0290	0.0277	0.0303	∞	0.0000	0.0000
μ_4	0.5114	0.5004	0.5229	∞	0.0000	0.0000
μ_5	1.0000			*not estimated*		
σ_1	0.0010			*not estimated*		
σ_2	0.0045	0.0042	0.0048	∞	0.0000	0.0000
σ_3	0.0249	0.0237	0.0261	∞	0.0000	0.0000
σ_4	0.3364	0.3295	0.3436	∞	0.0000	0.0000
σ_5	0.0010			*not estimated*		
c_1	-0.6959	-0.9773	-0.4082	∞	0.0006	0.0012
c_2	-0.0349	-0.3203	0.2510	1.4857	0.0006	0.0012
c_3	0.8952	0.6087	1.1777	∞	0.0006	0.0012
c_4	2.7509	2.4649	3.0421	∞	0.0007	0.0012
ζ_{EAD}	-0.1099	-0.1357	-0.0824	∞	0.0001	0.0001
ζ_{Facility}	0.2038	0.1495	0.2584	∞	0.0001	0.0001
$\zeta_{\text{Protection}}$	-0.4147	-0.4751	-0.3559	∞	0.0001	0.0001
ζ_{Industry}	-0.2355	-0.3000	-0.1683	∞	0.0002	0.0002
ζ_{HPI}	0.0590	-0.2183	0.3311	2.0243	0.0006	0.0010
			random effect			
σ	0.8191	0.6200	1.0329	∞	0.0005	0.0005

Note: The table summarizes the results of the LGD model. Parameters are stated in the first column. Categorical variables are included via dummy coding. The reference categories are term loan for facility, no for protection, and non FIRE for industry. The second column presents the posterior means. In the third and fourth column, lower and upper bounds of the corresponding HPDIs to a credibility level of 95% are displayed. The fifth column contains the posterior odds. Naive and time series standard errors are shown in the last two columns. Time series standard errors are calculated based on the effective chain length (N^*_{MCMC}) instead of the actual chain length (N_{MCMC}), whereby, $N^*_{\text{MCMC}} < N_{\text{MCMC}}$ holds for autocorrelated chains.

and have rather small standard deviations ($\sigma_2 = 0.0045$ and $\sigma_3 = 0.0249$), whereas, the forth component seems to cover the range in between the extremes of no and total loss ($\mu_4 = 0.5114$ and $\sigma_4 = 0.3364$). The posterior distributions of the cut points (c_k for $k \in \{1, 2, 3, 4\}$) are not directly interpretable as they depend on the range of the latent variable (Y^*).

Component probabilities are derived based on the OL model within the LGD model. The parameter of the EAD (ζ_{EAD}) exhibits a negative posterior mean indicating a lower value of the latent variable (Y^*) for higher EADs and, thus, lower LGDs. This impact is characterized by decisive evidence as the posterior odds are tending to infinity (posterior $\text{odds}_{\text{E}[\zeta_{\text{EAD}}]<0} \to \infty$) and the HPDI ($\text{HPDI}_{\zeta_{\text{EAD}}} = [-0.14, -0.08]$) excludes zero. Reasons for the negative impact of the EAD might be found in higher resolution efforts and, thus, lower loss rates, for loans of major size. The posterior mean of lines (ζ_{Facility}) is positive. Thus, lines are characterized by higher LGDs compared to term loans. This positive influence is decisively evident (posterior odds$_{E[\zeta_{\text{Facility}}]>0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\zeta_{\text{Facility}}} = [0.15, 0.26]$). Protection ($\zeta_{\text{Protection}}$) exhibits a negative posterior mean with decisive evidence (posterior odds$_{E[\zeta_{\text{Protection}}]<0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\zeta_{\text{Protection}}} = [-0.48, -0.36]$). This indicates lower loss rates for protected loans which corresponds to the economic intuition. According to the negative sign of the industry FIRE (ζ_{Industry}), LGDs for loans of this industry affiliation are lower compared to other industries. This impact is decisively evident (posterior odds$_{E[\zeta_{\text{Industry}}]<0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\zeta_{\text{Industry}}} = [-0.30, -0.17]$). The applied macro variable, i.e., the HPI (ζ_{HPI}), exhibits a positive sign indicating higher LGDs for higher values of the HPI. This contradicts the economic intuition as a sound economic surrounding should be accompanied with lower loss rates. However, the positive sign is not statistical evident (posterior odds$_{E[\zeta_{\text{HPI}}]>0} = 2.02 < 3.2$ and $0 \in \text{HPDI}_{\zeta_{\text{HPI}}} = [-0.22, 0.33]$). The last row of the table summarizes the posterior distribution of the random effect parameter.

Figure 5 illustrates the realizations of the random effect f_t in the LGD model. Higher realizations of the random effect ($f_t > 0$) indicate higher values of the latent variable Y^* for all loans defaulted in t and, thus, higher average LGDs in this quarter. The left panel of the figure presents the time

Figure 5: Random effect of the LGD model

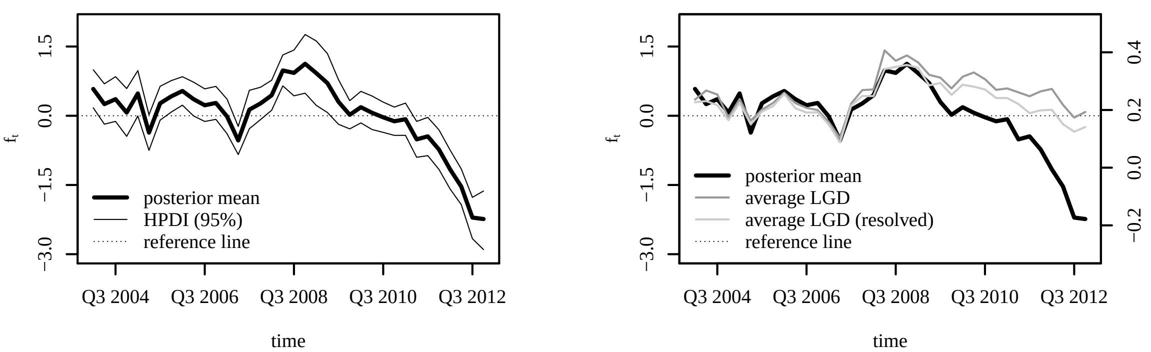

Note: The figure illustrates the course of the random effect in the LGD model over time. In the left panel, the posterior means (thick line) and the HPDI (95%, thin lines) of the random effect realizations, i.e., f_t, are displayed. In the right panel, the random effect (black line) is contrasted with the time patterns of average LGDs for all loans (dark gray line) and for resolved loans (light gray line). Final and incurrent LGDs as of validation sample I are included in the averaging. The dotted lines mark zero and serves as a reference line.

patterns of f_t. The path of f_t seems to be related to the economic cycle. While the realizations of the random effect scatter around zero prior to the crisis, increased values occur since 2007 Q2. In the climax of the GFC, f_t reaches its maximum. The rebound in the aftermath of the crisis instates gradually. The right panel of the figure contrasts these time patterns of f_t to average

LGDs in the time line as of Figure 4. The latter include observations which are not considered in the estimation, i.e., final LGDs of unresolved loans as of validation sample I. Up to the more recent time periods, the random effect seems to mimic the path of average LGDs. The time series disperse afterwards, whereby, the spread further increases in the time line. This deviation might be attributed to the neglection of censored observations. The final realizations of censored observations tend to have higher LGDs. This distorts the estimated realizations of the random effect in the more recent time periods. The effect of censoring and the associated distortion worsen in the time line, i.e., the distortion of f_t enlarges for higher t. Furthermore, a distortion of the random effect parameter σ has to be considered as the downward distortions in the random effect realizations might erroneously increase the underlying standard deviation of the random effect. We will come back to this later on (see subsequent paragraph).

Hierarchical Model

In analogy to the LGD model, the hierarchical approach is estimated based on the estimation sample (see Table 2). Due to DRT model in the hierarchical approach, it is possible to include censored observations, i.e., unresolved loans, in the estimation process. By this means, we are able to generate posterior predictive distributions for the DRT of unresolved cases and, thus, posterior predictive distributions for the LGD of unresolved loans. Furthermore, effects of the resolution bias as in the pure LGD model (see Figure 5) might be diminished.

Table 4 summarizes the results of the hierarchical model. Parameters are stated in the first column, whereas, the second column presents posterior means. Posterior distributions for the estimated component parameters (μ_k and σ_k for $k \in \{2,3,4\}$) and loan specific covariate parameters of the LGD model in the hierarchical approach (γ_{EAD}, γ_{Facility}, $\gamma_{\text{Protection}}$, and γ_{Industry}) are similar to their counterparts in the pure LGD model (see Table 3, μ_k and σ_k for $k \in \{2,3,4\}$ and ζ_{EAD}, ζ_{Facility}, $\zeta_{\text{Protection}}$, and ζ_{Industry}). A deviation arises for the parameter of the HPI (γ_{HPI}). In comparison the corresponding parameter in the pure LGD model (ζ_{HPI}) it exhibits an intuitively negative sign, thus, indicating lower LGDs in sound economic surroundings which is displayed by an increasing HPI. However, the parameter of the macro variable is still characterized by a lack of statistical evidence (posterior $\text{odds}_{E[\gamma_{\text{HPI}}]<0} = 1.18 < 3.2$ and $0 \in \text{HPDI}_{\gamma_{\text{HPI}}} = [-0.13, 0.12]$). The sign switch of γ_{HPI} compared to ζ_{HPI} might be due to the inclusion of the logarithmized DRT as explanatory variable in the LGD model of the hierarchical

Table 4: Results of the hierarchical model

	posterior mean	HPDI (95%)		posterior odds	naive standard error	time series standard error
	LGD model in the hierarchical approach					
μ_1	0.0000			*not estimated*		
μ_2	0.0064	0.0062	0.0067	∞	0.0000	0.0000
μ_3	0.0279	0.0268	0.0290	∞	0.0000	0.0000
μ_4	0.5033	0.4923	0.5144	∞	0.0000	0.0000
μ_5	1.0000			*not estimated*		
σ_1	0.0010			*not estimated*		
σ_2	0.0043	0.0040	0.0045	∞	0.0000	0.0000
σ_3	0.0234	0.0223	0.0244	∞	0.0000	0.0000
σ_4	0.3384	0.3314	0.3453	∞	0.0000	0.0000
σ_5	0.0010			*not estimated*		
c_1	-1.4391	-1.5803	-1.3004	∞	0.0003	0.0005
c_2	-0.5848	-0.7242	-0.4422	∞	0.0003	0.0006
c_3	0.5728	0.4306	0.7090	∞	0.0003	0.0005
c_4	2.6716	2.5262	2.8169	∞	0.0003	0.0005
γ_{EAD}	-0.1952	-0.2233	-0.1667	∞	0.0001	0.0001
γ_{Facility}	0.3259	0.2700	0.3840	∞	0.0001	0.0001
$\gamma_{\text{Protection}}$	-0.6291	-0.6932	-0.5676	∞	0.0001	0.0002
γ_{Industry}	-0.2736	-0.3437	-0.2036	∞	0.0002	0.0002
γ_{HPI}	-0.0061	-0.1287	0.1170	1.1847	0.0003	0.0005
γ_T	0.9996	0.9711	1.0280	∞	0.0001	0.0001
	DRT model in the hierarchical approach					
β_0	0.7341	0.6112	0.8521	∞	0.0003	0.0006
β_{EAD}	0.0512	0.0343	0.0678	∞	0.0000	0.0000
β_{Facility}	-0.0903	-0.1238	-0.0555	∞	0.0001	0.0001
$\beta_{\text{Protection}}$	0.1345	0.0981	0.1718	∞	0.0001	0.0001
β_{Industry}	-0.1555	-0.1954	-0.1141	∞	0.0001	0.0001
β_{VIX}	0.2731	0.1514	0.3946	7141.8571	0.0003	0.0004
s	0.8488	0.8395	0.8583	∞	0.0000	0.0000
	random effect					
σ_T	0.3424	0.2627	0.4327	∞	0.0002	0.0002
σ_L	0.3615	0.2696	0.4634	∞	0.0002	0.0003
$\omega_{T,L}$	0.1863	-0.1398	0.5031	6.3057	0.0007	0.0008

Note: The table summarizes the results of the hierarchical model. Parameters are stated in the first column. Categorical variables are included via dummy coding. The reference categories are term loan for facility, no for protection, and non FIRE for industry. The second column presents the posterior means. In the third and fourth column, lower and upper bounds of the corresponding HPDIs to a credibility level of 95% are displayed. The fifth column contains the posterior odds. Naive and time series standard errors are shown in the last two columns. Time series standard errors are calculated based on the effective chain length (N^*_{MCMC}) instead of the actual chain length (N_{MCMC}), whereby, $N^*_{\text{MCMC}} < N_{\text{MCMC}}$ holds for autocorrelated chains.

approach (γ_T) as further systematic variables, i.e., the VIX and the random effect of the DRT model, enter the LGD model through the DRT. The posterior mean of γ_T has a positive sign indicating higher LGDs for loans with higher DRTs. In Section 4.1 (see Figure 3), we determined this relation descriptively. The impact of the DRT is decisively evident (posterior $\text{odds}_{E[\gamma_T]>0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\gamma_T} = [0.97, 1.03]$).

In the DRT model of the hierarchical approach, loan specific covariates and a macro variable, i.e., the VIX, are included. The posterior mean of the EAD (β_{EAD}) exhibits a positive sign. Thus, loans of major size are accompanied with longer DRTs. This supports the thesis we stated in the previous paragraph. Financial institutions might undertake higher resolution efforts for loans of major size. This might increase the DRTs and simultaneously lower LGDs. Decisive evidence can be stated for the positive impact of the EAD in the DRT model (posterior $\text{odds}_{E[\beta_{\text{EAD}}]>0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\beta_{\text{EAD}}} = [0.03, 0.07]$). According to the negative posterior mean of lines (β_{Facility}), this facility type is accompanied with shorter DRTs compared to term loans. This impact is decisively evident (posterior $\text{odds}_{E[\beta_{\text{Facility}}]<0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\beta_{\text{Facility}}} = [-0.12, -0.06]$). In analogy to the EAD, the impact of facility is opposite in the LGD and DRT model of the hierarchical approach. While lines are characterized by shorter DRTs, they result in higher LGDs. Reasons may be found in divergent resolution efforts related to the size of the loan and its protection. The posterior mean of protection ($\beta_{\text{Protection}}$) exhibits a positive, decisively evident (posterior $\text{odds}_{E[\beta_{\text{Protection}}]<0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\beta_{\text{Protection}}} = [-0.69, -0.57]$), sign indicating longer DRTs for protected loans. The impact of protection is divergent among the models in the hierarchical approach ($\gamma_{\text{Protection}} < 0$ and $\beta_{\text{Protection}} > 0$). This might be due to the nature of protection itself. If loans are secured either by collateral or guarantees, efforts have to be taken to realize the protection value. This might extent DRTs, however, reduce LGDs when the protection value is realized. The industry affiliation FIRE (β_{Industry}) reveals a negative posterior mean, thus, it is connected to shorter DRTs. The sign is decisively evident (posterior $\text{odds}_{E[\beta_{\text{Industry}}]<0} \rightarrow \infty$ and $0 \notin \text{HPDI}_{\beta_{\text{Industry}}} = [-0.69, -0.57]$) and corresponds to the sign of the LGD model in the hierarchical approach ($\gamma_{\text{Industry}} < 0$ and $\beta_{\text{Industry}} < 0$). Resolution prospects in the FIRE industry might be limited compared to other industries due to less tangible assets. Thus, DRTs are short and LGDs low. To control for the impact of the macro economy, the VIX (β_{VIX}) is included in the DRT model of the hierarchical approach. Its posterior mean is positive and decisively evident (posterior $\text{odds}_{E[\beta_{\text{VIX}}]>0} = 7,141.86 > 100$ and $0 \notin \text{HPDI}_{\beta_{\text{VIX}}} = [0.15, 0.39]$). This entails longer DRTs in bad economic surroundings which corresponds to the economic intuition.

The parameters of the multivariate random effect as of Equation (11) are stated in the lower panel of Table 4. As the DRT is included in the LGD model of the hierarchical approach, the random effect of the DRT model (F_t^T) enters the LGD model. Thus, the aggregated systematic impact of the random effects on LGDs ($\mathcal{F}_t$) is the linear combination of $\gamma_T F_t^T$ and F_t^L :

$$\mathcal{F}_t = \gamma_T F_t^T + F_t^L \tag{15}$$
$$\sigma_{\mathcal{F}}^2 = \gamma_T^2 \sigma_T^2 + \sigma_L^2 + 2\gamma_T \sigma_T \sigma_L \omega_{T,L},$$

whereby, $\sigma_{\mathcal{F}}^2$ is the variance of the aggregated systematic effect. Considering the results of Table 4, the standard deviation $\sigma_{\mathcal{F}}$ of $\mathcal{F}_t$ amounts to 0.54. This standard deviation is considerably smaller compared to the standard deviation of the random effect in the pure LGD model (see Table 3, $\sigma = 0.82$). As suspected in the previous paragraph, the estimated standard deviation of the random effect in the pure LGD model seems to be distorted due to the resolution bias. Neglecting censored observations, i.e., unresolved loans, leads to distorted realizations of the random effect (f_t) and, thus, subsequently to distorted parameters (σ).

Figure 6 illustrates the realizations of the random effects of the DRT model f_t^T (upper left panel) and the LGD model f_t^L (lower left panel) in the hierarchical approach. Higher realizations of the random effect in the DRT model ($f_t^T > 0$) imply higher DRT for all loans defaulted in t, whereas, higher realizations of the random effect in the LGD model ($f_t^L > 0$) lead to higher values of the latent variable $\mathcal{Y}^*$ for all loans defaulted in t and, thus, to higher average LGDs in this quarter. Hence, DRTs impact LGDs through two channels (see Section 3). Directly, as higher DRTs are inserted in the LGD model. Indirectly, as positive realizations of f_t^T tend to imply positive realizations of f_t^L due to the positive correlation ($\omega_{T,L}$). However, the indirect channel might also weaken the impact of DRTs on LGDs as negative realizations of f_t^L are still possible. Considering the time patterns of the random effects as of Figure 6, four settings of the indirect channel are apparent. In the first setting prior to the GFC, $f_t^T < 0$ and $f_t^L > 0$ are valid. Thus, average DRTs of loans defaulted in t are shorter. The positive realization of f_t^L, however, increases average LGDs. Resolutions of these loans at least partly take place during the crisis. This might depress recovery payments at the end of the resolution process and, thus, increases LGDs. The second setting in the climax of the GFC is characterized by positive realizations of

Figure 6: Random effect of the hierarchical model

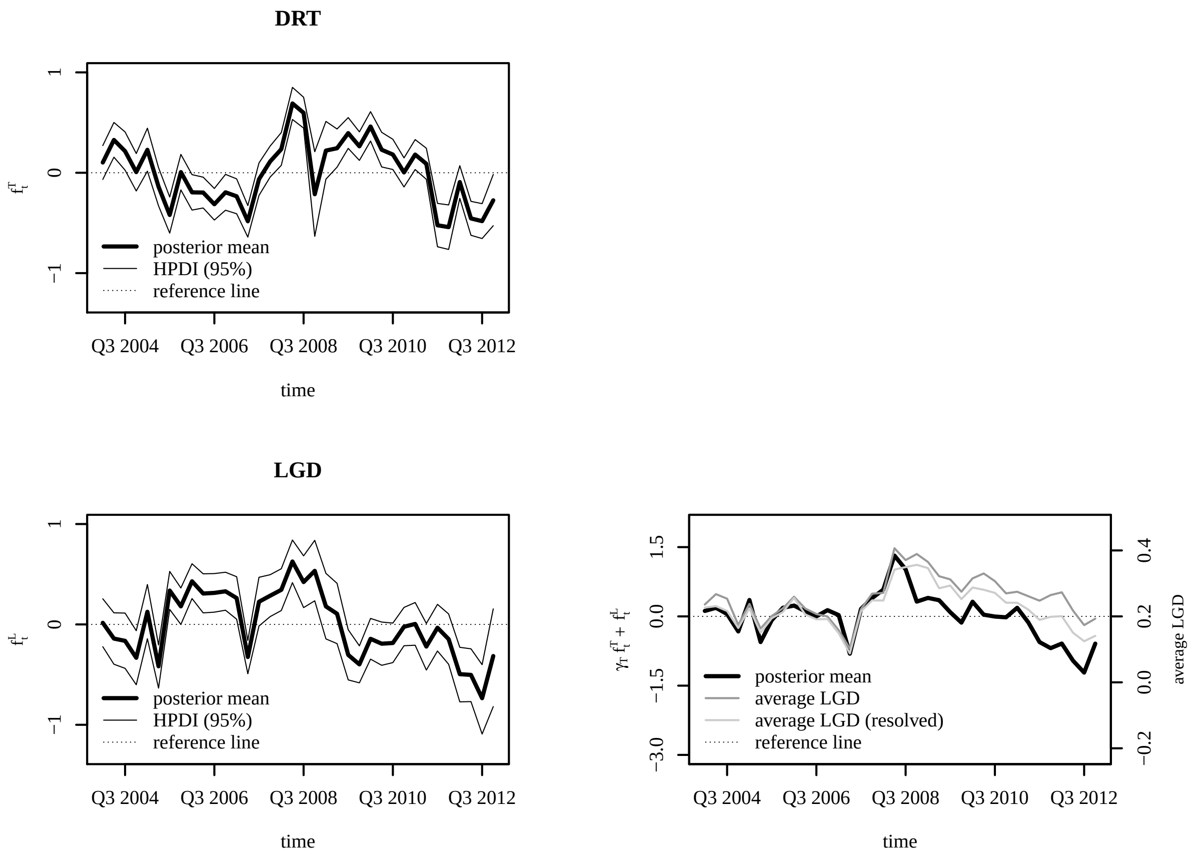

Note: Note: The figure illustrates the course of the random effects in the hierarchical model over time. In the left panels, the posterior means (thick lines) and the HPDI (95%, thin lines) of the random effect realizations, i.e., f_t^T (DRT) and f_t^L (LGD), are displayed. In the right panel, the combined systematic effect on the LGDs according to the random effects of the hierarchical model ($\gamma_T f_t^T + f_t^L$, black line) is contrasted with the time patterns of average LGDs for all loans (dark gray line) and for resolved loans (light gray line). Final and incurrent LGDs as of validation sample I are included in the averaging. The dotted lines mark zero and serves as a reference line.

both random effects ($f_t^T > 0$ and $f_t^L > 0$) indicating longer DRT and simultaneously higher LGDs of loans defaulted in t. In the third setting in the aftermath of the GFC, signs of the random effects are contrary ($f_t^T > 0$ and $f_t^L < 0$). Hence, average DRTs of loans defaulted in t are longer, whereas, average LGDs are lower. This might be due to the time delay as of the first setting. Analogously, parts of the recovery payments take place during the rebound period which favors recovery collection and decreases LGDs. The fourth setting is located in the most recent time period. The realizations of both random effects exhibit negative signs ($f_t^T < 0$ and $f_t^L < 0$) indicating shorter DRTs and simultaneously lower LGDs for loans defaulted in t. These settings illustrate the impacts of systematic effects in the resolution process. The positive correlation of the random effects ($\omega_{T,L}$) seems to be driven by extreme economic surroundings as synchronism appears in crises and boom periods. Furthermore, reasoning for the gradual rebound in the aftermath of the GFC can be provided (see Figure 4). While the random effect of the LGD model f_t^L indicates the rebound in the aftermath of the crisis (third setting), the random effect of the DRT model f_t^T remains on its high level. This might be due to the high stock of non-performing loans in the aftermath of the GFC which decelerated resolution proceedings. Average LGDs increase due to the direct channel.

The right panel of Figure 6 contrasts the aggregated systematic impact of the random effects ($\mathcal{F}_t$) to average LGDs in the time line. The latter include observations which are not considered in the estimation. The aggregated systematic effect seems to mimic the path of average LGDs. However, slight dispersions are apparent in the more recent time periods. Reasons might be found in a less accurate estimation of the random effect realizations of the LGD model (f_t^L) in the more recent time periods. Although censored observations are included through the DRT model, unresolved loans do not directly enter the LGD model in the hierarchical approach. Comparing the dispersions of the hierarchical model with the pure LGD model (see Figure 5), improvements are apparent. While the spread extremely increases in the time line for the pure LGD model, the deviation is considerably less pronounced in the hierarchical approach. Thus, the hierarchical approach succeeds in reducing distortions regarding the random effect.

5 Validation

As stated in Section 4.1 (see Table 2), the models are estimated based on the estimation sample. In the *in sample* validation, the posterior predictive distributions based on the estimation sample are compared to the empirical distributions of completely resolved loans in the estimation sample. The *out of sample* validation examines the distributional fit for censored observations, i.e., loans which have defaulted till the end of the estimation period but are still unresolved. Thus, Posterior predictive distributions based on validation sample I are compared to the corresponding empirical distribution. The posterior predictive distributions are generated based on the estimated realizations of the random effect. In the *out of sample out of time* validation, loans which defaulted after the end of the estimation period are considered. As no random effect realizations are available for those loans, posterior predictive distributions are generated on the means of the random effects, i.e., zero, and compared to the corresponding empirical distribution.

We adapt two graphical tools to evaluate the distributional fit of the models. First, kernel density estimates of the posterior predictive distributions are compared to kernel density estimates of the empirical data. The bandwidth is fixed to 0.015 to ensure graphical comparability. So heights of the kernel density estimates are comparable despite ties. Second, quantile-quantile plots are applied. Hereby, the quantiles of the posterior predictive distributions are contrasted to the quantiles of the empirical distribution. In the case of optimality, i.e., if the distributions correspond to each other, the points of the quantile-quantile plot are on the bisector line. If the probability of low loss components is overestimated and the probability of high loss components is underestimated, the points are below the bisector line as the the quantiles of the posterior predictive distributions are smaller than the quantiles of the empirical distribution. This corresponds to an underestimation of average LGDs.

In sample

Figure 7 illustrates the in sample validation of the LGD model and the hierarchical model. In the left panel, kernel density estimates of the empirical distribution (thin black line), the posterior predictive distribution of the LGD model (thick gray line), and the posterior predictive distribution of the hierarchical model (thick black line) are presented. However, lines lie directly on top of each other, thus, the black line of the posterior predictive distribution of the hierarchical model overlays the remaining two. To get a more detailed impression, the right panel of the figure illustrates quantile-quantile plots, whereby, the quantiles of the posterior predictive distributions are contrasted to the quantiles of the empirical distribution. The gray dots mark the LGD model, whereas, the hierarchical model is represented by black dots. The dots are near the optimality, i.e., bisector, line for both models. Thus, the in sample fit of the posterior predictive distributions is quite good with respect to the LGD model and the hierarchical model. This indicates that the applied FMM with five mixture components and two fixed components seems to deliver satisfactory results regarding the distributional fit.

Out of sample

Figure 8 illustrates the out of sample validation for the LGD model and the hierarchical model. The presentation corresponds to Figure 7 (in sample validation). The posterior predictive distribution of the LGD model (gray) seems to overestimate the probability mass of the low loss components, whereas, it underestimates the probability mass of the high loss components. In

Figure 7: Validation (*in sample*)

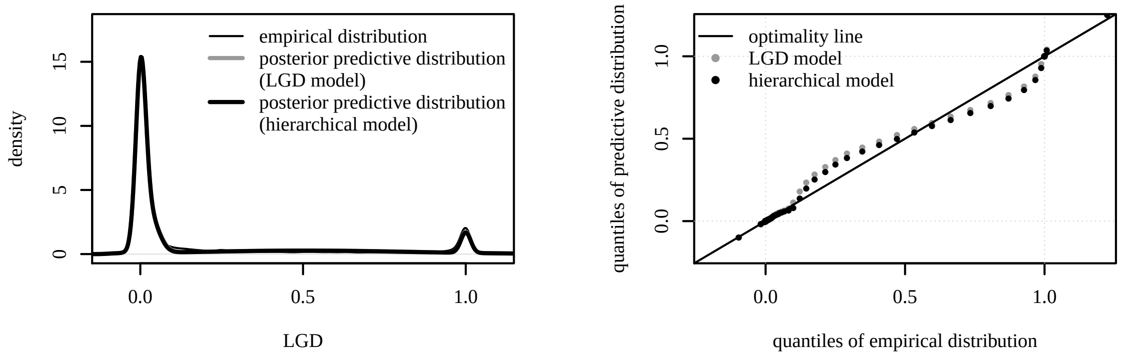

Note: The figure illustrates the *in sample* validation of the LGD model and the hierarchical model. In the left panel, the kernel density estimates of the empirical distribution (resolved loans as of estimation sample, thin black line), the posterior predictive distribution of the LGD model (thick gray line), and the posterior predictive distribution of the hierarchical model (thick black line) are displayed. The band width is fixed to 0.015 to ensure comparability. The kernel density estimates lie directly on top of each other. Thus, differences are not identifiable. In the right panel, the corresponding quantile-quantile plots are presented. The quantiles of the empirical distribution (x-axis) are plotted against the quantiles of the posterior predictive distributions (y-axis). The gray (black) dots mark the quantiles of the empirical distribution vs. the quantiles of the posterior predictive distribution of the LGD model (hierarchical model). The black line represents optimality.

the left panel, its kernel density estimate lies above the kernel density estimate of the empirical distribution for no loss (LGD = 0) and below for total loss (LGD = 1). This appears even clearer considering the quantile-quantile plots in the right panel. The dots are considerably below the optimality line indicating an underestimation of average LGDs. The underestimation is caused by the parameter distortion of the random effect due to the neglection of censored observations. The random effect realizations f_t are characterized by a downward bias, thus, leading to downward biased estimates for Y_i and downward biased loss rates (see Section 4.2). In contrast, the distributional fit of the hierarchical model is good on an out of sample perspective. This is due to two reasons. First, the parameter distortions are diminished (see Section 4.2). Second, the additional information of the time in default is utilized to improve predictions on an out of sample perspective. Final DRTs for censored observations, i.e., unresolved cases, are estimated within the hierarchical approach. These can be applied to generate predictions of final LGDs for unresolved loans.

Out of sample out of time

Figure 9 illustrates the out of sample out of time validation of the LGD model and the hierarchical model. The presentation corresponds to Figure 7. In analogy to the out of sample validation,

Figure 8: Validation (*out of sample*)

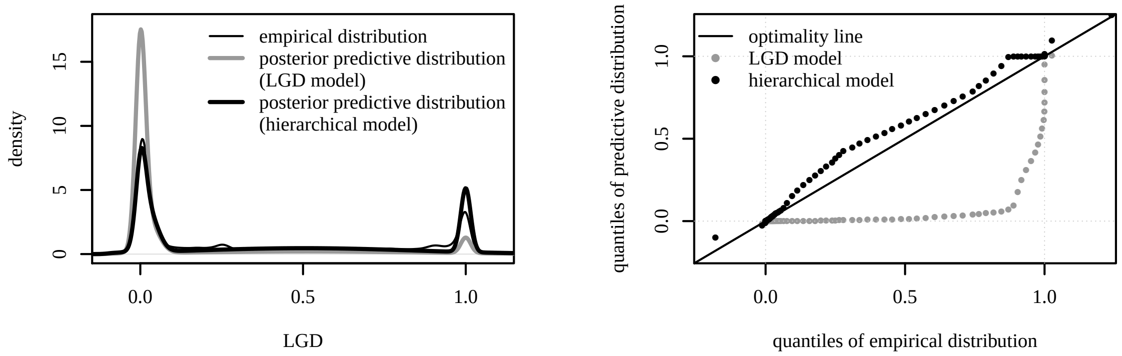

Note: The figure illustrates the *out of sample* validation of the LGD model and the hierarchical model. In the left panel, the kernel density estimates of the empirical distribution (resolved loans as of validation sample I, thin black line), the posterior predictive distribution of the LGD model (for resolved loans, thick gray line), and the posterior predictive distribution of the hierarchical model (for resolved loans, thick black line) are displayed. The band width is fixed to 0.015 to ensure comparability. In the right panel, the corresponding quantile-quantile plots are presented. The quantiles of the empirical distribution (x-axis) are plotted against the quantiles of the posterior predictive distributions (y-axis). The gray (black) dots mark the quantiles of the empirical distribution vs. the quantiles of the posterior predictive distribution of the LGD model (hierarchical model). The black line represents optimality.

an overestimation of low loss components and an underestimation of high loss components arises for the LGD model. However, it is not as striking as in the out of sample validation (see Figure 8). This might be due to the use of the random effect in average terms – instead of the individual realizations f_t as in the out of sample validation – to generate the posterior predictive distribution. However, the poor distributional fit of the LGD model on the out of sample out of time perspective suggests that there are additional distortions beyond the realizations of the random effect and its standard deviation. These might be found in the cut points which represent the intercepts in an OL model. The cut points of the LGD model and the hierarchical model are not directly comparable as the logarithm of the DRT is included as additional variable. By this means, the mean of the latent variable $\mathcal{Y}^*$ and, thereby, the level of the cut points are shifted. In contrast to the LGD model, the distributional fit of the hierarchical model is quite good on an out of sample out of time perspective.

Validation in the time line

Thus far, the distributional fit of the LGD model and the hierarchical model are analyzed for the estimation sample and the validation samples. Figure 10 illustrates the time patterns of average LGD predictions based on the posterior predictive distributions for specific default quarters. The upper left panel contrast average LGDs (thin black line) to average LGD predictions based

Figure 9: Validation (*out of sample out of time*)

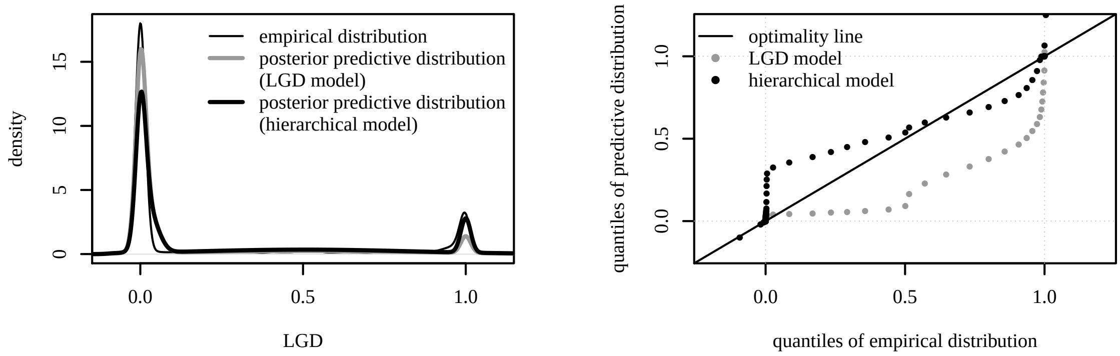

Note: The figure illustrates the *out of sample out of time* validation of the LGD model and the hierarchical model. In the left panel, the kernel density estimates of the empirical distribution (all loans as of validation sample II with incurrent LGDs, thin black line), the posterior predictive distribution of the LGD model (thick gray line), and the posterior predictive distribution of the hierarchical model (thick black line) are displayed. The band width is fixed to 0.015 to ensure comparability. In the right panel, the corresponding quantile-quantile plots are presented. The quantiles of the empirical distribution (x-axis) are plotted against the quantiles of the posterior predictive distributions (y-axis). The gray (black) dots mark the quantiles of the empirical distribution vs. the quantiles of the posterior predictive distribution of the LGD model (hierarchical model). The black line represents optimality.

on the LGD model (thick gray line) and the hierarchical model (thick black line) on an in sample perspective. In analogy to Figure 7, a good in sample fit for both models can be stated. The lower left panel illustrates the time patterns of average LGDs and LGD predictions on an out of sample perspective. Although the relative progressions of the LGD predictions based on the LGD model and the hierarchical model are similar, the predictions based on the LGD model are downward biased. Thus, average LGDs are underestimated by the LGD model in almost all quarters in validation sample I. This is not the case considering the predictions of the hierarchical model. The noisy behavior of average LGDs at the beginning of the time period is due to a lack of data as most loans defaulted in these quarters are resolved by the end of 2010 and, thus, not included in validation sample I. The lower right panel illustrates the time patterns of average LGDs and LGD predictions on an out of sample out of time perspective. The predictions based on the LGD model seem to be constant through time as the random effect is set to its mean, i.e., zero, and the macro variable is the only remaining systematic factor. However, the latter is not statistically evident (see Table 3). Furthermore, LGD predictions based on the LGD model seem to be systematically too low. LGD predictions based on the hierarchical model better fit average LGDs. Deviations at the end of the time period might be attributed to the inclusion of incurrent LGDs for unresolved cases (see Figure 4). Final LGDs will be lower and adjust the line downwards. In addition, LGD predictions based on the

Figure 10: Validation in the time line

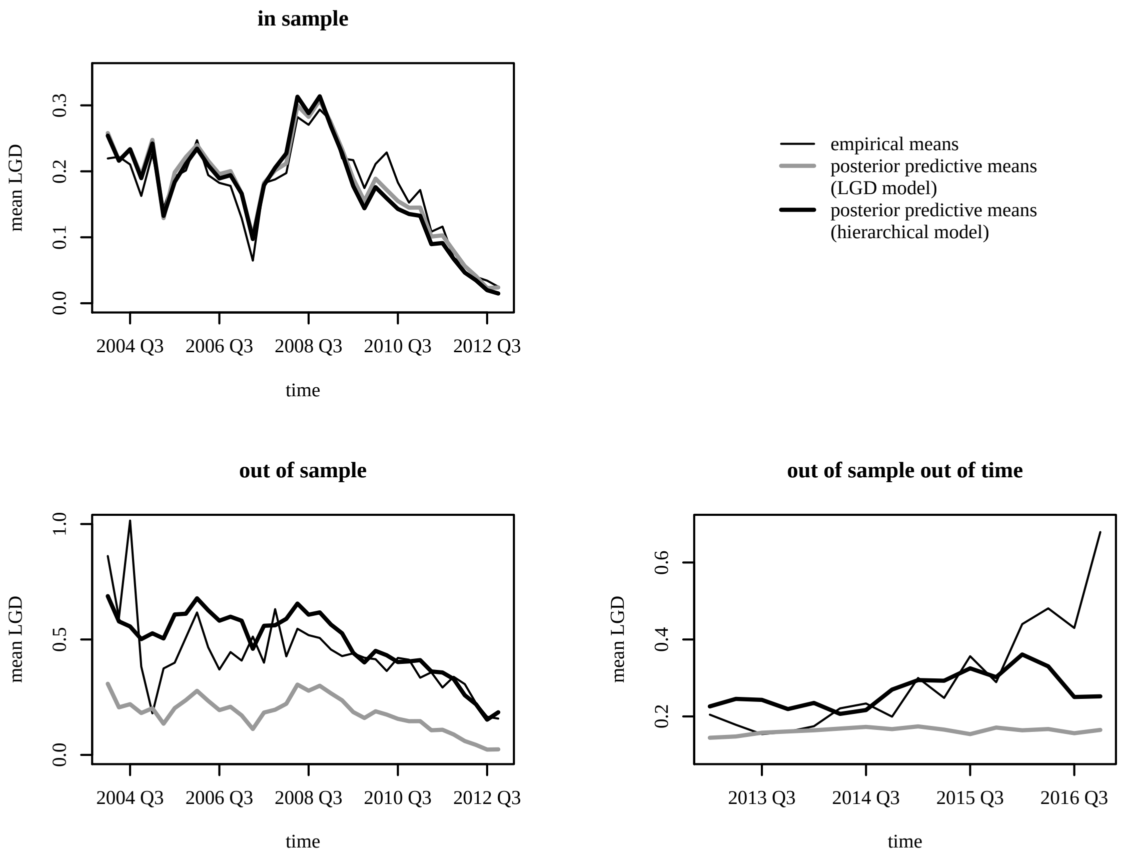

Note: The figure illustrates the validation in the time line. The means of the empirical distribution are displayed by a thin black line, whereas, the means of the posterior predictive distributions are marked by a thick gray line for the LGD model and a thick black line for the hierarchical model, respectively. In the upper panel, the *in sample* validation in the time line is presented (empirical means of resolved loans in estimation sample). The lower panels show the *out of sample* (empirical means of resolved loans in validation sample I) and *out of sample out of time* validation (empirical means of all loans in validation sample II with incurrent LGDs) in the time line.

hierarchal model display systematic movement as the statistically evident macro variable of the DRT model is enclosed in the LGD model of the hierarchical approach (see Table 4).

6 Conclusion

In this paper, we deeply examine the dependence structure of DRTs and LGDs using a hierarchical modeling framework. We find direct and indirect dependencies among the credit risk parameters. First, LGDs seem to be directly impacted by DRTs, i.e., longer resolution processes are accompanied with higher losses. Second, the parameters are characterized by common time patterns as correlation of the random effects in the individual models is positive. Due to the random nature of these effects, the dependence of DRTs and LGDs might be intensified or weakened in certain time periods. We find similar signs of the random effect realizations during the GFC and deviating signs pre and post crisis. Due to the consideration of direct dependency structures, we are able to generate intuitive LGD predictions for censored cases. These are of high practical relevance in the light of the recent EBA guidelines (see European Banking Authority, 2017). Besides LGD predictions for the non-defaulted exposure (unconditional predictions), financial institutions are demanded to predict LGDs for the defaulted exposure conditional on post-default information – such as the time in default. The hierarchical approach diminishes parameter distortions due to the neglection of censored observations in a pure (standard) LGD model and, thus, enables adequate unconditional LGD predictions for the non-defaulted exposure and consistent conditional LGD predictions for the defaulted exposure within one modeling framework.

In a conditional perspective (LGD predictions for defaulted exposure), a pure LGD model generates average LGD predictions underestimating actual average LGDs by up to 25 percentage points for loans defaulted during the GFC (2008 Q1 to 2009 Q3). The hierarchical approach delivers sufficiently conservative predictions for loans defaulted in the crisis (up to 16 percentage points above actual average LGDs). Assuming ten loans with an EAD of 500,000 EUR defaulted in 2008 Q4, a pure LGD model underestimates losses due to these loans by round about 1,05 million EUR. In a unconditional perspective (LGD predictions for non-defaulted exposure), these effects are less pronounced, however, still remarkable. A pure LGD model constantly underestimates actual average LGDs in the time period from 2013 Q1 to 2015 Q4 by up to 20 percentage points, while the hierarchical approach delivers slightly conservative predictions in most time periods (between 3 percentage points below and 8 percentage points above actual LGDs). Subsequent time periods are hard to interpret due to incurrent LGDs. Assuming ten loans with an EAD of 500,000 EUR defaulted in 2015 Q4, a pure LGD model underestimates losses due to these loans by round about 600,000 EUR.

Concluding, the consideration of censored observations is essential to generate suitable LGD predictions. The presented hierarchical model prevents the need of additional data constraints and provides fruitful insights into the dependence structure of DRTs and LGDs.

References

Altman, E. and E. Kalotay (2014). Ultimate recovery mixtures. *Journal of Banking and Finance 40*(1), 116–129.

Basel Committee on Banking Supervision (2004). International convergence of capital measurement and capital standards: A revised framework. *Bank for International Settlements*.

Basel Committee on Banking Supervision (2005). Guidance on paragraph 468 of the framework document. *Bank for International Settlements*.

Betz, J., R. Kellner, and D. Rösch (2016). What drives the time to resolution of defaulted bank loans? *Finance Research Letters 18*(1), 7–31.

Betz, J., R. Kellner, and D. Rösch (2018). Systematic effects among loss given defaults and their implications on downturn estimation. *European Journal of Operational Research 271*(3), 1113–1144.

Betz, J., S. Krüger, R. Kellner, and D. Rösch (2017). Macroeconomic effects and frailties in the resolution of non-performing loans. *Journal of Banking and Finance*, forthcoming.

Bijak, K. and L. Thomas (2015). Modelling LGD for unsecured retail loans using Bayesian methods. *Journal of the Operational Research Society 66*(2), 342–352.

Bris, A., I. Welch, and N. Zhu (2006). The costs of bankruptcy: Chapter 7 liquidation versus Chapter 11 reorganization. *Journal of Finance 61*(3), 1253–1303.

Calabrese, R. (2014). Downturn loss given default: Mixture distribution estimation. *European Journal of Operational Research 237*(1), 271–277.

Denis, D. and K. Rodgers (2007). Chapter 11: Duration, outcome, and postreorganization performance. *Journal of Financial and Quantitative Analysis 42*(1), 101–118.

Dermine, J. and C. Neto de Carvalho (2006). Bank loan losses-given-default: A case study. *Journal of Banking and Finance 30*(4), 1219–1243.

European Banking Authority (2017). Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposure.

Gelman, A. and D. B. Rubin (1992). Inference from iterative simulation using multiple sequences. *Statistical Science 7*(4), 457–472.

Gürtler, M. and M. Hibbeln (2013). Improvements in loss given default forecasts for bank loans. *Journal of Banking and Finance 37*(7), 2354–2366.

Heidelberger, P. and P. Welch (1981). A spectral method for confidence interval generation and run length control in simulations. *Communications of the ACM* (Special issue on *Simulation Modeling and Statistical Computing*) *24*(4), 233–245.

Heidelberger, P. and P. Welch (1983). Simulation run length control in the presence of an initial transient. *Operations Research 31*(6), 1109–1144.

Helwege, J. (1999). How long do junk bonds spend in default? *Journal of Finance 54*(1), 341–357.

Höcht, S., A. Kroneberg, and R. Zagst (2011). Explaining aggregated recovery rates. *Working paper*.

Höcht, S. and R. Zagst (2010). Loan recovery determinants – A Pan-European study. *Working paper*.

Kalotay, E. and E. Altman (2017). Intertemporal forecasts of defaulted bond recoveries and portfolio losses. *Review of Finance 21*(1), 433–463.

Kass, R. and A. Raftery (1995). Bayes factors. *Journal of the American Statistical Association 90*(430), 773–795.

Krüger, S. and D. Rösch (2017). Downturn LGD modeling using quantile regression. *Journal of Banking and Finance 79*(1), 42–56.

Loterman, G., I. Brown, D. Martens, C. Mues, and B. Baesens (2012). Benchmarking regression algorithms for loss given default modeling. *International Journal of Forecasting 28*(1), 161–170.

McLachlan, G. and D. Peel (2000). *Finite mixture models*. Wiley.

Partington, G., P. Russel, M. Stevenson, and V. Torbey (2001). Predicting return outcomes to shareholders from companies entering Chapter 11 bankruptcy. *Managerial Finance 27*(4), 78–96.

Qi, M. and X. Zhao (2011). Comparison of modeling methods for loss given default. *Journal of Banking and Finance 35*(11), 2842–2855.

Stan Development Team (2016). Stan modeling language. *User's Guide and Reference Manual*.

Wong, B., G. Partington, M. Stevenson, and V. Torbey (2007). Surviving Chapter 11 bankruptcies: Duration and payoff? *Abacus 43*(3), 363–387.