Improving PD and LGD models; following the changes in the market

Wemke van der Weij[a]*, Marcel den Hollander[a]

[a]Group Risk Management, SNS REAAL, Utrecht, The Netherlands

July 23, 2009

Abstract

We consider the modelling of the probability of default and the loss given default in the context of credit risk control for retail portfolios. The loss due to the non-payment of a debtor needs to be predicted so that financial institutions can hold sufficient capital to safeguard its solvency and overall economic stability. We first explain how a classical model works and then continuing with its main drawback we discuss calibration methods of the model. We give several directions for calibration where we take into account all available information with respect to the historical data. Next, using lifelike data, we illustrate the calibration methods by applying our method to a case study based on mortgages in The Netherlands.

1 Introduction

Credit risk, the loss due to a debtor's non-payment of a loan or other line of credit, plays an important role in financial institutions. And these financial institutions are strong players in the economic market. To control credit risk, guidelines are formulated in the so-called Basel II agreement. This agreement was initially published in 2004. With respect to credit risk the guidelines roughly imply that the greater risk to which a bank is exposed, the greater the amount of capital a bank needs to hold to safeguard its solvency and overall economic stability. An import measure for the prediction of the risk is Expected Loss (EL) which is derived from Loss Given Default (LGD), Probability of Default (PD), and Exposure At Default (EAD) estimators. By definition in the Basel II guidelines a borrower is in default if he is more than 90 days past due on principal or interest on any material obligation to the bank.

The LGD is based on the total loss on a default to infinity. The time horizon is bounded for practical reasons, considering materiality of cash flows in time.

*corresponding author: Wemke van der Weij, email: Wemke.vanderWeij@SNSREAAL.nl

We chose to use a 24 month performance period to finish the default and obtain all relevant cash flows. The PD is by definition the probability of a borrower getting into the default status within 12 months. Note that in practice the 24 months time frame for the LGD is not sufficient to cover all relevant, since this assumes that there are no cash flows after 24 months which is not often not the case. The requirement of a defined time for observation of the default and related cash flows introduces an appealing complexity due to the definition of the estimators, especially with a rapid movement in the economy. Due to this performance period a structural lag exists between the estimators and their realisations. The PD, LGD and EAD estimators are given as the outcome of a model which is based on historic data of at least three years, according to Basel II.

In the literature there are several papers on modelling the PD and LGD estimators and the recovery rates which are closely related to PD and LGD estimators. Grunert and Weber [4] analyse the recovery rates using data of German companies and find that several variables are significant for predicting the recovery rates. Their main result is the observation that intense client relationship leads to higher recovery rates. In [2] recovery rates and default rates are linked to macroeconomic variables and added to a model based on historical data. Also Acharya et al. [1] are interested in the relations of macroeconomic variables to the default and recovery rates linking distress scenarios of an industry to recovery rates. In [3] an econometric model is proposed in which the unobserved Markov chain is introduced to represent the so-called credit cycle. In this model data of the past in used instead of the covariates of the previous mentioned models. In contrast to the literature we assume a given model for estimation the credit risk and focus on calibration methods to improve the prediction. Our method overcomes the complexity of fitting the estimators given their realisation, which occurs due to the time lag between these two.

Our solution relies on taking into account fractions of the realisations instead of a realisation in a fixed performance period. These fractions can serve as historical input for the model to predict the estimators even before the performance period ends. These fractions are scaled to adjust the previous observed model outcomes. This results in more agile estimators following the volatility of the economy. The proposed approach significantly improves the punctuality of the estimators and thus of the Expected Loss. To illustrate the methodology and the results we present a case study based on mortgages in The Netherlands. The setup of the paper is as follows. In Chapter 2 the starting point, the basic model, for our approach is explained. Next, in Chapter 3 new calibration methods are considered. And Chapter 4 illustrates these methods with a case study. Conclusions and topics for further research are listed in Chapter 5.

2 The model

In this section we discuss the model which is the basis for the calibration methods discussed in the next section. However, if the underlying model for predicting the credit risk is slightly different than the model discussed in this paper the calibration methods in the next section is probably still applicable. In the Basel II framework, borrowers are assigned to distinct classes based on their risk characteristics. And in each class the clients have more or less equivalent risk profiles. The minimum number of classes for the PD and LGD estimators is defined in the Basel II guidelines. A common way to categorise the clients in risk classes is by using historical data to identify and quantify the client characteristics and external variables. These explanatory variables are called the risk factors.

We consider the following situation; separate models for PD and LGD estimation, using regression techniques on given datasets. For both estimators, the models are scorecards based on the most relevant variables. The relevant variables are selected both via evaluation of the statistical performance and expert judgement of key representatives of relevant departments. This ensures that the variables in the models are fit to quantify risk in an explainable manner. The variables reflect characteristics of the client, his behaviour, the product and the relevant securities. Common examples are respectively age, amount of late payments and the last two combined, the relative size of a loan to the value of the underlying security.

Risk classes are defined as ranges of scores. Each client is put in a risk class by evaluating the individual score. The individual scores are the sum of the scores given to the explanatory variables of clients. Consider for example a client receiving 30 points for his age and 100 points for his behaviour. Based on his total score of 130 he is placed in a risk class. Given a risk class (a bucket) a prediction of the estimated value is added. This estimated value is based on realisations in the past. Additionally sometimes economic variables such as employment rate or gross domestic product are included in the model. However, it is very difficult to specify the futures impact of such [2].

The estimated PD and LGD values can be calibrated. This calibration can also include the macroeconomic cycle to some extend. Several restrictions on the models must still be valid after calibration. The most relevant is that the buckets are significantly different and monotonically increasing over the buckets (sorted by score). This ensures a valid distinction between the clients. This is also the main requirement when using other methods to categorise clients.

The expected loss calculations for Basel II also require an EAD estimation. The EAD is calculated as follows: the outstanding loan minus the savings at the bank and plus the outstanding debt. This approach is chosen since netting is allowed in our scenario and the loan is relatively large compared to monthly mutations.

3 Calibration

For calibration we only focus on the PD and LGD estimators, since the EAD estimation is a very straight forward calculation performed on each individual client. As explained in the previous section the borrowers are placed in buckets for PD and LGD separately and for each bucket the LGD and PD estimators are predicted. While the explanatory variables for the model can still be valid, resulting in a classification of the clients in the risk buckets, the values of the LGD and PD estimators for each bucket might be outdated. Initially these values are estimated concurrently with the model development and are based on historical data. However, due to the common assumption that the PD and LGD values can only be predicted based on borrowers getting into default 12 (PD) or respectively 24 (LGD) months ahead a time lag exists between realisations and predictions. A simple calibration of the model, including recent data, does not overcome this problem, but only slightly corrects for the gap between the estimators and the realisations. Therefore we propose a new method that uses the recent data more efficiently and effectively. Our method can be applied on all levels of the original model, for instance on bucket and on portfolio level. If applied on portfolio level an additional step is needed to translate the calibrated figures to bucket level.

First we recap a very common calibration method, which we call the level-calibration. At the introduction of the model the estimated values were predicted using a historical time frame. After some time the value can be calibrated by estimating the values based on realisations in a more recent period. However, this more recent time frame will still be related to events that took place quite some time ago due to the time-lag between realisations and estimators. Thus, even when a simple level-calibration is applied it is always far behind the macro economic situation. Therefore we propose a method based on fractions of the realisations so that we can include the most recent realisations instead of only the realisations of an LGD and a PD of respectively 24 and 12 months ahead and overcome the complexity of the time lag.

We first introduce general notation, which we will later specify for the LGD and PD model. Denote by y_t the estimator at moment t and let z_t be the realisation of this estimator. Our aim is to find the calibrated estimator, $\hat{y}_t$, where this estimator should be better than the previous estimator y_t. To validate the statement 'better' we first introduce three different measures which are commonly used for the purpose as described above. We compare the Mean Square Error (MSE), the Mean Absolute Error (MAE) and the Mean Absolute Percentage Error (MAPE) of y_t and $\hat{y}_t$. The errors are defined as follows:

MSE

$$\sqrt{\frac{1}{2}\sum_{t=1}^{n}(\hat{y}_t - z_t)^2} < \sqrt{\frac{1}{2}\sum_{t=1}^{n}(y_t - z_t)^2}, \tag{1}$$

MAE

$$\sqrt{\frac{1}{2}\sum_{t=1}^{n}|\hat{y}_t - z_t|} < \sqrt{\frac{1}{2}\sum_{t=1}^{n}|y_t - z_t|^2}, \tag{2}$$

MAPE

$$\sqrt{\frac{1}{2}\sum_{t=1}^{n}|\frac{\hat{y}_t - z_t}{z_t}|} < \sqrt{\frac{1}{2}\sum_{t=1}^{n}|\frac{y_t - z_t}{z_t}|^2}. \tag{3}$$

At the moment y_t is known the realisation of this estimator is not known, and least the performance period of the estimator, which we denote by T. Note that given a time instance τ we can only compare the estimator of time $\tau - T$ with its realisations. For more recent estimators the realisations are still in the future.

Now define a realisation fraction $z_t(m)$ as the realisations at time t given the action takes place in the m^{th} month. And let

$$z_t = \sum_{m=0}^{T} z_t(m). \tag{4}$$

For the default rate a fraction $z_t(1)$ equals the number of clients at time t getting in default after 1 month defined by the Total number of clients at moment t. For the realised loss rate the fraction $z_t(1)$ denotes the loss in the first month after the estimation period t.

Month Period	1	2	$\dots$	T
1	$z_1(1)$	$z_1(2)$	$\dots$	$z_1(T)$
2	$z_2(1)$	$z_2(2)$	$\dots$	$z_2(T)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
$t-2$	$z_{t-2}(1)$	$z_{t-2}(2)$	$-$	$-$
$t-1$	$z_{t-1}(1)$	$-$	$\dots$	$-$
t	$-$	$-$	$\dots$	$-$

Table 1: Realisation fractions.

The above description of realisation fractions will be translated to LGD and PD estimators in the next paragraph.

Let us first consider the LGD. The LGD has a performance period of 24 months ($T = 24$). Denote by $\text{RLR}_t(m)$ the Realised Loss Rate (the realisation of the LGD) given the borrowers getting in default at period t and resulting in a loss in the m^{th} month after default. The RLR_t equals the Loss (L) divided by the EAD, where again the index t can be added to denote the loss or EAD respectively for a default period of length t:

$$RLR_t = \frac{L_t}{EAD_t}. \tag{5}$$

We obtain the following relation:

$$RLR_t = \sum_{m=1}^{T} \frac{L_t(m)}{EAD_t(m)} = \sum_{m=1}^{T} RLR_t(m), \tag{6}$$

with t defined in months and $L_t(m)$ and $EAD_t(m)$ defined similarly to $z_t(m)$.

Similarly for the PD we introduce realisation fractions, namely the default rate (DR) with index t (DR_t) which denote the fraction of borrowers getting in default in the t^{th} month.

Recap Table 1 and consider the PD, then the first column represents the default rate for clients getting into default in the first month. The next column represents the default rate of clients getting into default in the second month. Note that the last row is not filled, since given we are at time instance $t-1$ we can only observe the clients getting into default in the first month since the second and other months are still in the future. Thus each diagonal represents the most recent observable default rates given the time instance in the left column. The diagonal represents all borrowers getting into default in the month as stated by the most left column where the diagonal starts, the in-default cohort at time t of borrowers.

As stated before our approach is to use the most recent realisation fractions. Given Table 1 the most recent realisations are $z_{t-1}(1) \ldots z_{t-T}(T)$. Instead of using only these realisations, or only using complete realisation (a row in Table 1) of time $t-T$ our calibration methods use the more recent observations slightly different.

Note that if we use only one diagonal, $z_t(1) \ldots z_{t-T}(T)$, all the realisations take place in the same month. And this might be to sensitive to the period considered. Therefore we mention several more sophisticated manners to take into account the most recent realisation fractions. We consider the following three manners:

- Moving Average (MA),
- Linear Regression (LR), and
- Exponential Smoothing (ES).

The Moving Average (MA) is a rolling average. Instead of just summing the last diagonal one can take a moving average of length k. Depending on the length of the rolling period the new values will follow the realisations more accurate. The idea is to apply the MA as follows:

$$MA(k) = \hat{y}_t = \sum_{m=1}^{T} \left(\frac{\sum_{i=t-k}^{t} z_i(m)}{k} \right) \tag{7}$$

Linear regression can be done using the same realisations but instead of taking a rolling average a linear regression is performed estimating the trend and the level for each column of Table 1 of a length $t-k$. Last method we mention is the exponential smoothing which is almost similar to Equation 7, but exponential weights assigns exponentially decreasing weights as the observation get older.

The advantage of the linear regression, following a linear trend is unfortunately also the drawback of the method. Since the realisations might not follow a linear trend use of this method can lead in the wrong direction. Moving average uses historical data and is more prudent in that sense. However it always follows the correct trend of the data with a time lag depending on the length of the moving average. Exponential smoothing might improve the calibration but due to seasonal effects it can perform contrarily to what is preferred. Given these considerations in the next section we investigate in the moving average using data of Dutch mortgages.

Remark 3.1 *As stated before the cash flows (realisations related to the LGD) and even the end of default moment do not always fall within the performance period T. Therefore a common assumption is that after T the loss is given by a fraction of the exposure for open defaults, depending on the loans and the securities of these loans. Note that this loss, which is in fact an approximation of the real loss will be used as input for the calibration and is thus considered to be a real observation. If there are many borrowers reaching this period T the impact of this approximation on the estimators $\hat{y}$ is significant. Alternately it is possible to include data of defaults where cash flows are received after this performance period. But note that the number of observations will be very low, and using the realisation fractions can only lead to a better predictor if there are enough observations in each considered interval length (column/ k rows) combination as presented in Table 1.*

4 Case Study

In this section we apply the described method for the LGD model on a fictional dataset, based on mortgage data of the Dutch market. The data is very volatile, which is typical for LGD sets.

One complexity with calibration, especially in a downturn, is the possible postponement of foreclosures. This can result in selection effects where the lower graded cases do not lead to realisations. To anticipate on this we assume a fixed percentage return on a foreclosure sale after a certain time. This fictional return is used to derive the RLR value for such a client two years after start of the default. The RLR values that are obtained like this are used with the RLR values for the other cases which are based on the observed cash flows.

In the preceding figures we observe that the realisations are generally at a higher level than the estimations. Note that the realisations are known at a different moment than the estimates. The estimates are known after the monthly processes are finished. The realisations are only known in their final form after the performance period (T) from the period on the x-axis. In this example the realisations after March 2007 will alter when the last defaults are closed.

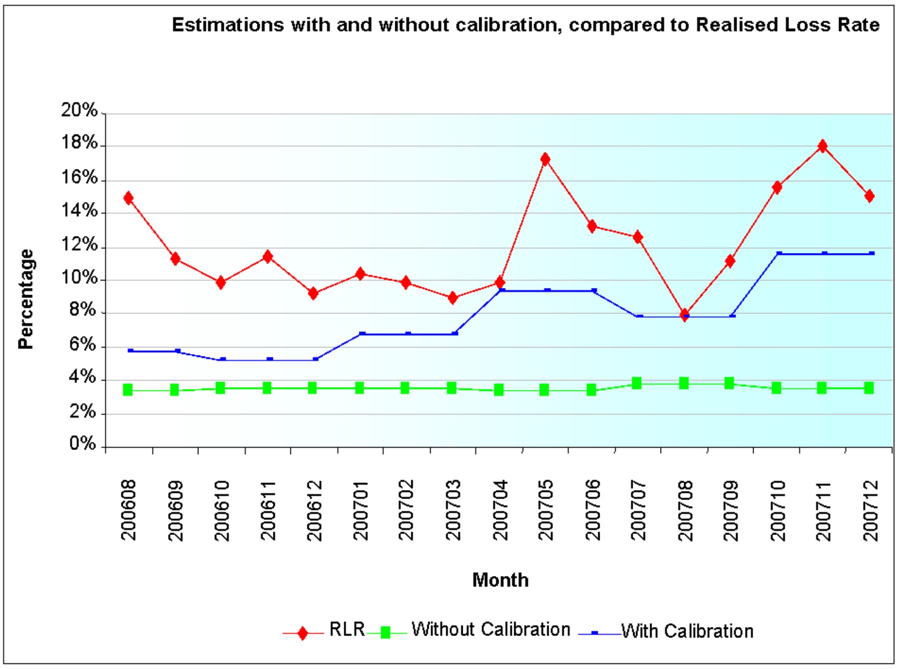

Figure 1: Comparison of the model with calibration using MA(12) and without calibration

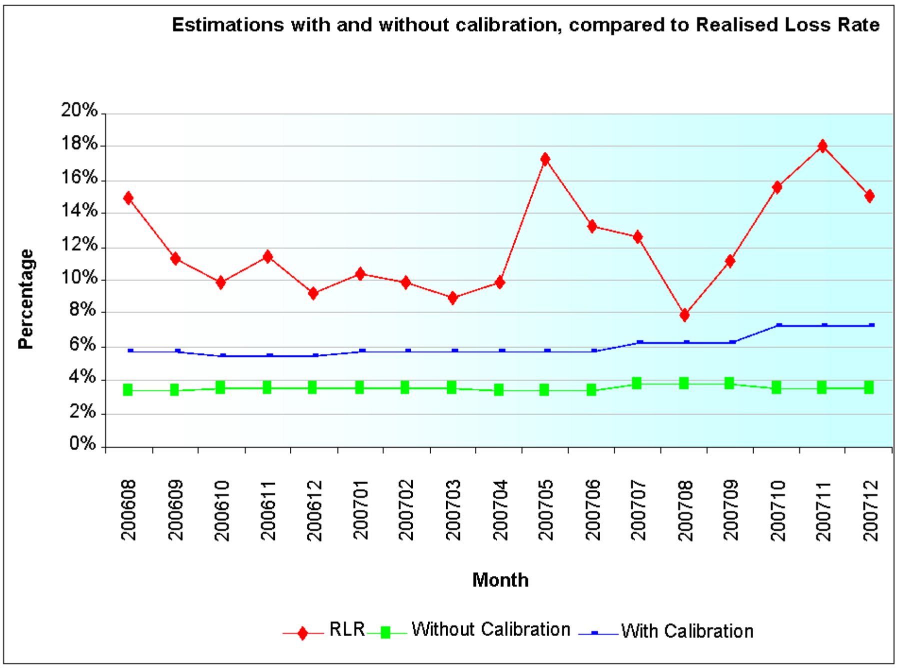

Figure 2: Comparison of the model with calibration using MA(1) and without calibration

Measure	Without	MA(12)	MA(1)
MSE	$10,17\%$	$7,85\%$	$5,42\%$
MAE	$9,64\%$	$7,30\%$	$4,90\%$
MAPE	$75,73\%$	$56,37\%$	$37,80\%$

Table 2: Comparison of the errors for estimators with and without calibration, calculated for the period January 2007 to December 2007 for MA (12).

The calibrated estimations are moving up with the realisations. In the method we do not make assumptions on future developments, so the method is basically just closing the gap with the realisations.

The improved statistical performance is illustrated in the following table:

The main considerations when selecting the best method are the statistical performance and the usability in conjunction with relevant business processes. The statistical measures favour the MA(1) model. For business processes, the preferred option is MA(12) over MA(1) because it is less sensitive to annual cycles and ensures more stability in reported figures. The economic cycle could be included in the estimation, but is beyond the scope of this case study.

5 Conclusion and further research

In this paper we consider the prediction of the expected loss and it's building blocks, which is very important for the estimation of the credit risk of financial institutions and it is required by the Basel II guidelines. We first discussed several methods and outline a model based on historical data, where the borrowers are scored in risk buckets using client characteristics such as behaviour and age. This very common model first scores the clients in buckets and then adds a value to that bucket. In the next section we proposed new ideas for calibration of this value, using the recent data more sophisticated reducing the gab between the realisations and the estimates.

This paper focuses mainly on calibration of the estimated values and is illustrated by a case study. In this paper the underlying model is not discussed. However, an opening to improve the performance of the model is to estimate the values not only once on the realisations but also on explanatory variables, which can include seasonal patterns, economic variables and variables representing the stability and changes in the processes in the organization. Also the calibration might possibly be improved further if the link between the economic cycles and process can be made with the estimators and the number of observations. And this number of observations can be used as input for the exponential smoothing method or for the number of the considered realisation fractions.

References

[1] V.V. Acharya, S.T. Bharath, and A. Srinivasan. Does industry-wide distress affect defaulted firms? Evidence from deritor recoveries. *Journal of Financial Economics*, 85:787–821, 2007.

[2] E.I. Altman, B. Brandy, A. Resti, and A. Sironi. The link between default and recovery rates: theory, emperical evidenxe, and implementations. *Journal of Business*, 78:2203–2227, 2005.

[3] M. Bruche and González-Aguado C. Recovery rates, default probabilities, and the credit cycle. *Journal of Banking & Finance*, 4(9), 2009.

[4] J. Grunert and M. Weber. Recovery rates of commercial lending: Emperical evidence for german companies. *Journal of Banking & Finance*, 33(3):505–513, 2008.