Macroeconomic conditions in models of Loss Given Default for retail credit

Tony Bellotti and Jonathan Crook
Credit Research Centre
University of Edinburgh Business School
William Robertson Building
50 George Square, Edinburgh EH8 9JY

Abstract

Loss Given Default is an important measure of credit loss used by financial institutions to compute risk within credit portfolios, expected loss on individual loans and capital requirements. We investigate models of Loss Given Default for UK retail credit cards which incorporate macroeconomic conditions. We find bank interest rates and unemployment rate are important explanatory variables and their inclusion improves forecasts of Loss Given Default for hold-out data sets. Additionally, the inclusion of macroeconomic conditions is important since it enables stress testing and provides a means to model downturn Loss Given Default as required by the Basel II Accord for the advanced internal ratings based approach to calculating capital requirements.

Key Words: Loss given default; retail credit; regression; model comparison; macroeconomic; forecast.

1. Introduction

Loss Given Default (LGD) is the loss incurred by a financial institution when an obligor defaults on a loan, given as the fraction of exposure at default (EAD) unpaid after some period of time. It is usual for LGD to have a value between 0 and 1 where 0 means the balance is fully recovered and 1 means total loss of EAD. LGD is an important value that banks need to estimate accurately for several reasons. Firstly, it can be used along with probability of default (PD) and EAD to estimate expected financial loss. Secondly, a forecast of LGD for an individual can help determine the collection policy to be used for that individual following default. For example, if high LGD is expected, then more effort may be employed to help reduce this loss. Thirdly, an estimate of LGD, and therefore portfolio financial risk, is an integral part of the operational calculation of capital requirements to cover credit loss during extreme economic conditions. The Basel II Capital Accord [2006] allows banks the opportunity to estimate LGD using their own models with the advanced internal ratings based (IRB) approach.

In this paper we focus on modelling and forecasting LGD for UK retail credit cards based on account variables (AVs) with the inclusion of macroeconomic variables (MVs). Our prior expectation is that as interest rates rise, so the cost of mortgage and other debt increases, making it more difficult for an obligor to repay outstanding credit card balances and so increasing the mean LGD. Equally, an increase in unemployment level means more people find themselves in circumstances where they cannot repay credit and this would also increase the mean LGD. On the other hand, an increase in earnings would mean more people have more income available to pay off debt and would therefore decrease mean LGD. In addition it is possible that some defaulters are more likely to be less able to repay than others when the state of the economy changes. For example, those who are unemployed at the time of credit card application may be particularly sensitive to interest rate increases, as may home owners. Similarly borrowers with a higher default balance may be particularly sensitive to increases in interest rates. For this reason we also consider interactions between MVs and account data.

Given that the distribution of LGD is a bimodal U-shape, we may consider a Tobit model or a decision tree model along with various transformations of the dependent variable. However, in a previous study, we found that the best forecasting model is Ordinary Least Squares (OLS) regression (Bellotti and Crook 2007). Therefore in this study we simply use OLS regression. Economic conditions are included as values of MVs for bank interest rate, unemployment level and earnings growth at time of account default. We find that the first two MVs are statistically significant explanatory variables that give rise to improved forecasts of LGD in hold-out tests at both an account and portfolio level. Building LGD models with MVs also addresses the Basel II requirement to estimate “downturn LGD” since stressed values of MVs can be used in the model to forecast LGD during poor economic conditions. We demonstrate how this can be done by stressing interest rate values.

The modelling and forecast of LGD for retail credit using macroeconomic conditions is a new area of study. There is an extensive literature regarding LGD models for corporate loans (see eg Altman et al [2005]). However, there is less about *forecasting* LGD. An exception is Gupton and Stein [2005] who describe a predictive LGD model for corporate loans using Moody-KMV’s Losscalc© software. There is also very little literature regarding retail credit LGD, even though this is a large financial market: total lending in the UK consumer credit market reached over £1.4 trillion in 2009 [source: Bank of England]. Grippa et al [2005] publish empirical LGD models for a sample of 20,724 Italian accounts that includes small businesses along with households. They observe differences in LGD and recovery periods across different geographic regions and different recovery channels. They also conducted a multivariate analysis that showed a statistically significant negative relationship between the presence of collateral or personal guarantee and LGD, and a positive relationship with size of loan. However, the range of variables used is far more limited than would be available to a financial institution that has made credit card or personal loans and the study did not attempt to forecast LGD. Dermine and de Carvalho [2005] model LGD for loans to small and medium sized firms in Portugal. They apply mortality analysis and include annual GDP growth as an explanatory variable. However, they found that GDP growth was not significant. They suggest that this may be due to the fact that the period of analysis, 1995 – 2005, did not include a significant recession. We may also note that their training sample size (374 defaults) was relatively small and may not have been large enough for a significant relationship between the economy and LGD to be discovered. Querci [2005] provides an LGD model for loans to small businesses and individuals by an Italian bank. This study shows the importance of regional differences on LGD variation but does not include time varying macroeconomic conditions. Figlewski et al [2007] model the effect of macroeconomic factors on corporate default with a detailed study of numerous economic conditions including unemployment level, inflation, GDP and a production index. They found that many of these MVs were significant explanatory variables.

The novelties of our paper are that unlike published work (1) we consider forecasts of LGD for retail credit cards, (2) we report results of model comparison, (3) we include macroeconomic conditions in our models and (4) we do this using a very large sample across several different credit card products. In section 2 we describe our modelling and performance assessment methods. In section 3 we discuss the application and macroeconomic data used. In section 4 we provide model comparisons and test results along with a description of an explanatory model with MVs. Finally in Section 5 we provide some conclusions and discussion.

2. Method

We consider models with different combinations of variables and discuss our hold-out test procedure.

2.1 Models

In general, for retail credit, there are five categories of circumstances that will affect the amount an individual repays on a defaulted loan and can be used to build models of LGD:

(1) individual details, some of which can be collected at time of application such as age, income, employment, housing status and address;

(2) account information at default: date or age of account at default and outstanding balance;

(3) changes in personal circumstances of an obligor over time;

(4) macroeconomic or business conditions on date of default, or possibly with a lag or lead on date of default;

(5) operational decisions made by the bank, such as the level of risk they were willing to accept on the credit product and the process they use to follow up bad debt.

Of these, the richest source of explanatory variables we have is the information provided at time of application for credit along with the credit bureau score collected by the bank at time of application. This is data falling into category (1). We also include category (2) data, account information at default. Including this data implies the model is conditional on default. It is possible to build models unconditionally but this is outside the scope of this paper. It is difficult for a lender to extract data in category (3). It cannot easily keep track of an individual's employment status or, even less so, his or her personal difficulties, such as divorce or illness, that may lead him or her to be unable to fully repay debt. It is possible to use account behaviour data or a behavioural score but we do not do this in this study since such information is not homogeneous within the data we have. It is understood that LGD is likely to be time dependent, varying over the business cycle [Schuermann 2005]. Therefore we include macroeconomic conditions (4). Including a bank's operational decisions (5) for each credit card product over time could also be fruitful. However, this information was not available for our study.

We build and compare models with and without MVs. We also explore models that include interaction terms between AVs and MVs. We restrict our models to OLS since our previous study showed this was sufficient for LGD forecasting (Bellotti and Crook 2007). The four models we consider are:-

1. **Simple**: no covariates in the model. This model effectively forecasts the same mean LGD taken from the training data for all test cases.
2. **AV**: account variables only.
3. **AV&MV**: account and macroeconomic variables.
4. **AV&MV with interactions**: also includes interaction terms between AVs and MVs. It is not feasible to include all interaction terms so variable selection is used as described later in Section 3.3.

As is conventional in the literature we model LGD in terms of *recovery rate* (RR) rather than LGD directly, where RR = 1-LGD. Our working definition of RR is

$$\mathrm{RR} = \frac{\text{sum of repayments made over a period } t \text{ following default}}{\text{outstanding balance at date of default}}.$$

The choice of recovery period is a business decision. We consider a recovery period of t=12 months following default. Often banks are interested in longer periods or want to estimate recovery at close of account, but we found that a 12 month model can be used to estimate for longer recovery periods. It is possible to calculate LGD in alternative ways. For example, time after default could be an event, such as account charge off, rather than a fixed period; or the market value of the bad debt based on sale of the exposure in the market could be taken into account along with repayments; or administration costs of following up default may be included in the LGD calculation. However, these alternatives are beyond the scope of this paper. Loss of interest payments could also be included in the definition of LGD but this is not required by Basel II.

2.2 Model assessment

For OLS we report *adjusted* R^2 for model fit. This is for two reasons. Firstly, the various models we consider are nested so inevitably those with additional covariates will give improved R^2 model fit. The adjusted R^2 compensates for the additional variables. Secondly, reporting R^2 is misleading since it does not give a fair comparison between different studies with different sample sizes.

We report coefficient estimates for the OLS model with RR as dependent variable. However, since RR is not normally distributed, the error terms may not be normally distributed. Therefore conventional estimators of standard error may be biased. Instead we use a bootstrap to construct distributions for the coefficient estimates [Kennedy 2003, section 4.6]. As Lam and Veall [2002] show when OLS is used with non-normal, and in particular bimodal, distributed error terms, the bootstrap gives accurate estimates of confidence intervals where the usual analytic method fails.

We report the relative effect of each MV within the model. The coefficient estimates for MVs are multiplied by their standard deviation in the training data to derive *standardized coefficient estimates*. They show change in RR for a one standard deviation change in the covariate value and are therefore comparable and give an indication of the relative importance of each MV within the model. This approach was taken to study the effects of MVs on corporate default by Figlewski et al [2007]. Since our credit card data spans the period from 1999 to 2005, standard deviations for each MV are computed based on values within this period.

To test the effectiveness of the LGD model for forecasts, we use a hold-out sample, testing on credit card default data that is independent and follows chronologically after the period of the training data used to build the models. This approach allows us to simulate the expected operational use of LGD models in retail credit when a financial institution may want to assess the LGD risk on a new batch of defaults based on the performance of past defaults. In detail, we select cohorts of test data sets consisting only of accounts that default in a particular quarter. For each of these cohorts we train only on default data available prior to that quarter. Since we need to measure LGD after a period of t months following default, we need to ensure that date of defaults in the training data are at least t months prior to the beginning of the test quarter. This test procedure is illustrated in Figure 3. So, for example, if our test set was 2003Q3 and we are considering LGD after 12 months, our training data would consist of all cases that default within the period from 1999Q1 to 2002Q2. To estimate robust models with MVs we should train over the whole business cycle which is usually considered between 3 to 5 years. Therefore we consider taking training data sets with a minimum of 3 years of defaults. This then gives us 10 quarters of test set data from 2003Q1 to 2005Q2. Results across these independent test sets then form a time series of forecast results.

The accuracy of forecasts relative to the observed true values is measured at the account level by mean square error (MSE). However, we are also interested in how well the model is able to estimate the observed, or true, mean LGD or RR over a portfolio of accounts. Therefore, for each test set quarter we measure the difference between forecast and observed mean RR across all test cases. If this difference is greater than zero, then the model is generally overestimating RR, whereas when it is less than zero, it is underestimating RR. The closer the difference between forecast and observed RR is to zero, the better the estimate. The mean value of both the MSE and absolute value of difference in forecast and observed mean RR (abs diff RR) across the several test quarters are reported in order to get an aggregate measure of performance.

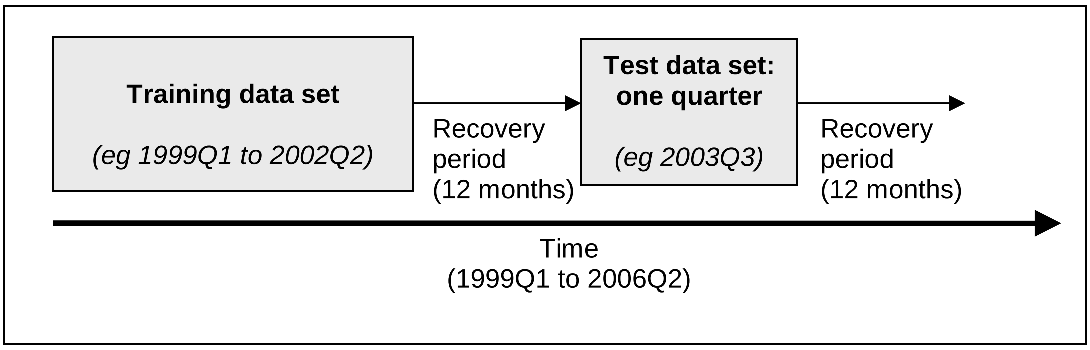

Figure 3. Illustration of hold out test procedure for one quarter of test data.

3. Data

3.1 Application data

For this study we have available a data set consisting of over 55,000 credit card accounts in default over the period 1999 to 2005 for customers across the whole of the UK. Account holders are expected to make a minimum payment of outstanding balance each month. We define *default* as a case where a credit card holder has failed to make minimum payments for three consecutive months or more. This is a typical definition of default for credit cards [Thomas et al 2002, p.123] and is in line with the default definition given by the Basel II Accord [BCBS 2006]. The data consists of four different credit card products which are a selection from those offered by a financial institution. As is typical for LGD, its distribution in our data set is between 0 and 1 and approximately U-shaped[1]. The calculation of LGD should ideally include administration costs for managing and implementing a collection procedure following account default. Unfortunately, this information was not available for this data set. The credit card data we use has many details extracted at time of application. These

[1] For reasons of commercial confidentiality, to protect our data supplier, we cannot reveal exact distributions or values of LGD for our data, nor descriptive details of other variables. This is usual in this research area when large portfolios of live financial data are being studied. Nevertheless since our focus is on reporting forecast performance, this should not be a hindrance to scientific reporting.

include the applicant's housing and employment status, age, income, total number of known credit cards and length of relationship with bank (time with bank). A credit bureau score is also provided and is from the same source and so is homogeneous across products. Additionally, information at time of default is also included in the data. This consists of balance outstanding at default and age of credit card account. We include balance at default since there is strong evidence from past studies that it is an important effect [Grippa et al 2005; Dermine and de Carvalho 2006] and it makes sense to include it operationally, especially if it improves forecasts of LGD. Some variables, time with bank, income and age, have a small percentage of missing values (less than 6% of accounts). For each of these variables, we code missing values to 0 and create a dummy variable to indicate a missing value to capture the mean value amongst accounts with missing values.

Several studies have found that PD and LGD are positively correlated (see Rösche and Scheule (2006)). We also found this to be the case in our data. Nevertheless, PD is not included in our models since it is effectively represented by including the AVs that are usually used to model it, along with MVs that may explain the joint systemic risk to both PD and LGD (Altman et al 2005).

Any single credit card portfolio is liable to have operational effects that will alter overall risk over time for that specific product, such as changes in cut-offs on credit score when accepting applications. This may lead to idiosyncratic links to economic conditions and therefore poorer models using MVs. By combining data across several products the impact of these idiosyncratic effects will be reduced and changes in risk over time are more likely to be linked to more objective effects such as the economy. Additionally, combining several products into one data set will increase the training set size. These two factors should lead to stronger MV. A dummy variable is used to indicate which product the account belongs to, in order to model different levels of RR between products.

3.2 Macroeconomic variables

We consider these three series of macroeconomic data for the UK which we believe would have a strong direct effect on mean LGD for UK retail credit cards:-

- Selected UK retail banks' base interest rates.

- UK unemployment level: measured as thousands of adults (16+) unemployed.
- (Earnings) UK earnings index (2000 = 100) for the whole economy including bonuses as a ratio of the retail price index.

These are all available from the UK Office for National Statistics as monthly data. We use non-seasonally adjusted data for earnings since we expect that seasonal changes in the economy may have some effect on abilities to repay. We would also have preferred to use non-seasonally adjusted data for the unemployment level but unfortunately this was not available. GDP growth for the UK is a common indicator of economic conditions but we have not included it since it is not available as monthly data which is the granularity that lenders typically require and therefore the granularity we require for our models. The MVs are included for each case at the date of default but it may be that if there is a relationship between MVs and LGD then this is lagged or led. For example, changes in interest rates may, in general, affect ability to pay several months later. However, we found that changing lag or lead period does not lead to improved results.

Each MV has a time trend: interest rates and unemployment level are generally falling over the period 1999-2005 whilst real earnings are steadily increasing. Indeed, earnings generally increase exponentially with time therefore we include it in our model as growth in log earnings over 12 months to remove this obvious time trend. The three MV time series we use are shown in Figure 4. We ensure that any model fit of MVs as explanatory variables is not simply because they follow a time trend that matches a trend in RR by including date of default explicitly in the models. If a MV is a good explanatory variable simply because of a time trend, then the inclusion of date of default should weaken its effect within the model. Including date of default in the AV model also allows us to test whether improvements in forecasts are simply due to a time trend rather than specifically economic conditions.

High correlation between MVs is a potential problem since this could lead to multicollinearity within the LGD model and therefore distort parameter estimates. We can test for multicollinearity by measuring the variance inflation factor (VIF) given by $\left(1-R^2\right)^{-1}$ when each MV is regressed on all other model covariates

(Kennedy 2003). A high VIF indicates multicollinearity and a VIF greater than 5 is an indication that there may be a problem.

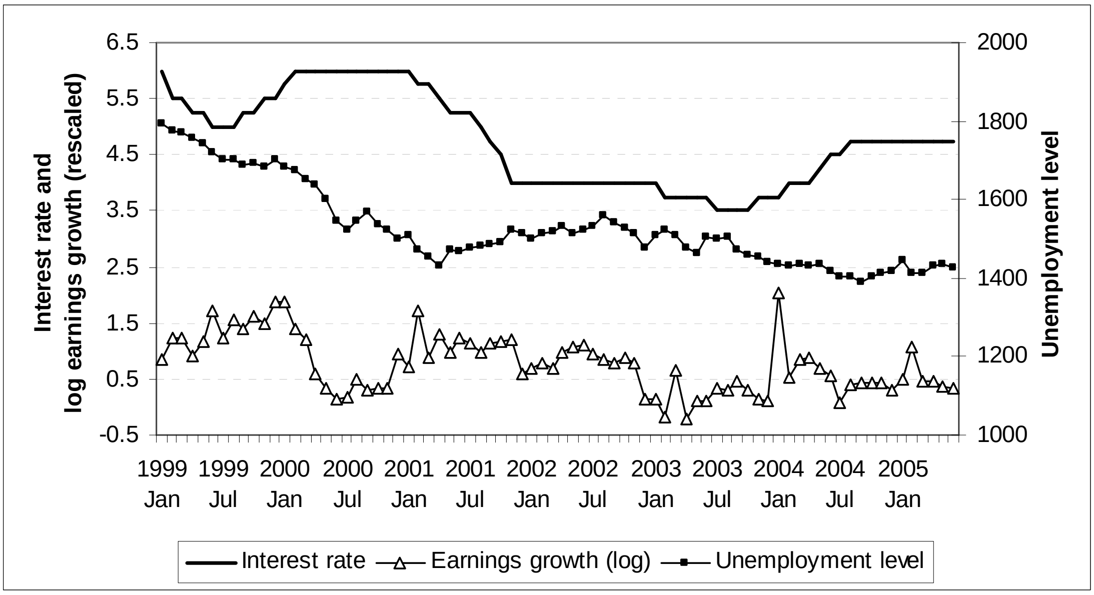

Figure 4. UK macroeconomic data 1999-2005.

3.3 Inclusion of interaction terms

Since there are many possible combinations of variables to form interaction terms, the number included in the model is controlled using forward variable selection. All AVs are included in the model but an iterative process is used to include MVs and interaction terms between AVs and MVs. At each step, each of the outstanding MVs and interaction terms not already in the model are added separately. The term that maximally increases a fit criterion is added to the model. The process is repeated until no new interaction terms are found that improve fit. There are several possible fit criteria that could be used and it is common to use an F-test. However, since we are interested in forecasting, we use Akaike's information criterion (AIC) [Akaike 1973]. This has the advantage that it takes account of the parameter space of the model and discourages complex models with large numbers of variables. In turn, this discourages over-fitting to the training data set. We approximate AIC by $n\ln(\mathrm{MSE})+2p$ where n and p are number of observations and number of parameters in the model respectively and MSE is the mean square error for observations in the training data. We find using the AIC criterion gives better forecasts than using the standard F-test. Further discussion of variable selection methods and use of AIC for predictive models is given by Miller [1990]. The variable selection procedure we use is further constrained so that for each interaction term included, its constitutive terms are also automatically included [Brambor 2005].

4 Results

Section 4.1 describes forecast performance for comparison of the different models and section 4.2 describes the best performing model for forecasts and its statistically significant explanatory variables. Stress tests are considered in sections 4.3.

4.1 Model comparisons

Table 1 shows forecast results for different models. It is clear that the model with both AVs & MVs performs best for both measures of forecast performance. The inclusion of interaction terms gives slightly worse forecasts than the AV&MV model.

Table 1. Aggregate forecast results for different models. Abs RR diff is the absolute difference between observed and forecast mean RR over each test quarter. Results are given as mean values across the series of test quarters from 2003Q1 to 2005Q2 or 2004Q2 for 12 or 24 month recovery periods respectively. Figures in bold show best results for each recovery period and product group.

Model		**Result**	
	Explanatory variables	**MSE**	**Abs RR diff**
1.	None (simple)	0.168	0.0391
2.	AV	0.156	0.0640
3.	AV & MV	**0.151**	**0.0148**
4.	AV & MV & interaction terms	0.152	0.0252

Figures 5 and 6 show forecast time series results for models 1 to 3 listed in Table 1 for each test quarter. They show that the AV&MV model performs consistently better over time. Figure 5 shows that MSE is lower for AV&MV than either the simple or AV model and improves over time, possibly a result of having training data over a longer period of the business cycle to model the MV effects. Figure 6 shows that overall the AV&MV model forecasts mean RR much more closely than the other models as is evidenced by how close the difference between forecast and observed mean RR is to zero. In contrast the AV model is consistently over-estimating RR over time. In 2004Q2, the simple model achieves the best result but this is serendipitous since its forecasts are simply moving from underestimating to overestimating RR at that time. These results show that including MVs is important to improve LGD forecast results. It should be noted, however, that in our study at least 3 years of training data is required. We found that using less than this led to unstable models with MVs that occasionally extremely over or under estimate LGD.

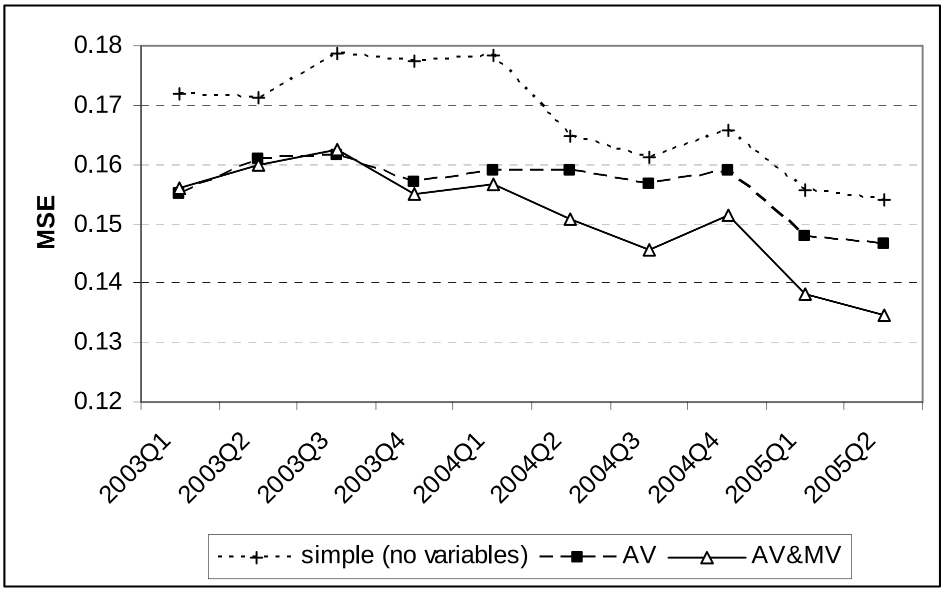

Figure 5. MSE of forecasts for each test data set from 2003Q1 to 2005Q2 for three OLS models with different variables.

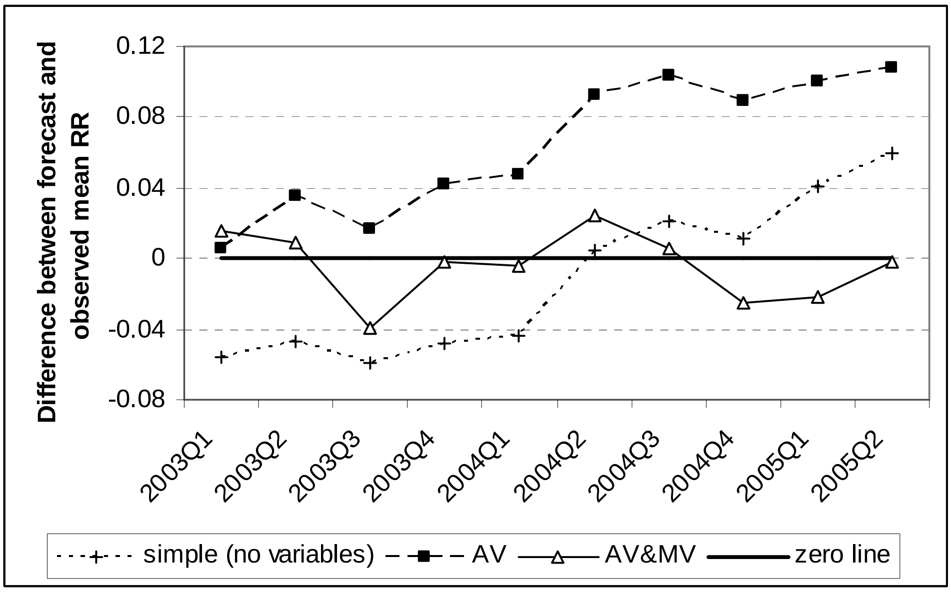

Figure 6. Difference between forecast and observed mean RR. Results relate to each test data set from 2003Q1 to 2005Q2 for three OLS model with different variables.

4.2 Explanatory model

Table 1 shows that the AV&MV model was the best performing forecaster so we describe this model estimate in further detail. We report model fit results for LGD models built on all data from 1999 to 2005. The AV model has an adjusted R^2 model fit of 0.105. When MVs are included this increases to 0.110. When interaction terms between MVs and application variables are also included using forward selection, the adjusted R^2 is 0.111. This small increase indicates that adding interactions does not give a noticeable improvement reinforcing the results observed for forecasting. Table 2 shows coefficient estimates for MVs in the AV&MV model. Coefficient estimates for AVs are not shown since they are not the focus of this study, but the estimate for the time trend in defaults (ie date of default) is shown since that is relevant to our interpretation of MVs. Since the error residuals are non-normal we have used bootstrap to compute statistical significance. The reported p-values are from a normal-based distribution imposed on bootstrap coefficient estimates.

Table 2. Coefficent estimates for date of default and MVs in model of recovery rates. Training data set for the AV&MV model included all available cases from 1999 to 2005. Standard errors and p-values are computed with bootstrap estimation with 1,000 repetitions. Values in *italics* show standardized estimates for MVs.

Variable	Coefficient estimate	Standard error	z	P>|z|
Date of default	-0.000010	0.000006	-1.62	0.105
Bank interest rates *standardized estimate*	-0.0508 *-0.0584*	0.00294	-17.3	0
Unemployment level *standardized estimate*	-0.000150 *-0.0199*	0.000051	-2.96	0.003
Earnings growth (log) *standardized estimate*	0.717 *0.0028*	0.484	1.48	0.139

Table 2 shows that coefficient estimates for MVs have the expected signs. That is, the parameter estimate for interest rates is negative meaning that higher interest rates at the time of default tend to give lower RR. Similarly, higher unemployment levels are also linked to lower RR. However, higher earnings growth, year-on-year, leads to increased RR which suggests earnings growth leads to better recoveries. Bank interest rates and unemployment level are both statistically significant at a 0.01 level, although earnings growth is not statistically significant in the model. The coefficient estimate for date of default is small and not statistically significant which implies that the effect of the MVs is not due to a simple time trend. Date of default is a stronger explanatory variable in the AV model and is statistically significant at a 0.01 level. Nevertheless, the results in Table 1, comparing AV and AV&MV, show that the time trend is not sufficient to explain the improved performance given by the model with MVs. Table 2 also shows standardized parameter estimates for MVs. These give an indication of the relative size of effect of each MV. Bank interest rate clearly has the largest magnitude with unemployment level having less than half the interest rate effect. Additionally we find that the VIF for any of the MVs when regressed on all other covariates in model (1) was always less than 2 which is sufficiently small that we should not expect that the results are affected by multicollinearity. Since including interaction terms did not improve forecasts or model fit we do not report interaction terms in the explanatory model.

4.3 Stress test using interest rates

We can use LGD models with MVs for stress testing by replacing values of MVs by stressed values. Therefore they can be used to model particular scenarios or as part of a larger simulation based on a distribution of past macroeconomic conditions. We conduct a univariate stress test by considering the effect of interest rate changes on the test data set in 2005Q2 which is the last quarter of test data available from our data set. Figure 4 shows that for the training period, extreme values of interest rate are 6% and 3.5%. When these values are used in the AV&MV model, the forecast mean RR for the test quarter changes by -17% and +28% respectively. This shows the shift in overall RR for the worst and best cases. A downturn RR 17% less than during normal or good conditions is not implausible.

5. Conclusion

Our database spanned the period 1999 to 2005. Figure 4 shows that this period covered a range of economic conditions in the UK with interest rates generally decreasing and an overall reduction in unemployment. Earnings generally rose, although at some times growth was higher than others. This period does have the disadvantage for our analysis that there were no major recessions or downturns to train from and towards 2005 the UK economy was stable and fairly unremarkable. This is a point also noted by Dermine and de Carvalho [2005] with regard to their study. We speculate that a very good macroeconomic model of LGD should have training data across the entire business cycle. Unfortunately, due to practical reasons of data availability, we were unable to provide this. Nevertheless, given this limitation, we still found the MV model to be effective. We show that adding bank interest rates and unemployment level as MVs into a LGD model yields better model fit and that these variables are statistically significant explanatory variables. Additionally including these MVs improves forecasts with generally better MSE and estimates of mean RR across test quarters. Although the improved MSE is modest, Figure 5 suggests that the AV&MV models improve relatively with the duration or size of the training data set. Comparing the AV&MV model with the AV model in Figure 6 shows a clearly better forecast of LGD at the portfolio level.

We found that the inclusion of interaction terms between AVs and MVs did not generally improve performance and led to slightly worse results. The poor performance of the model with interaction terms affirms the comment by Gayler [2006] that the main effects are believed to be more stable than interactions for prediction in credit scoring. Nevertheless we feel that there is likely to be some useful interaction effects between MVs and application terms; eg those with high outstanding debt, say a mortgage on a property, are more likely to be affected by changes in bank interest rates. The problem is to determine which of them are important prior to modelling. The automated forward selection process we use is clearly insufficient for this task. Gayler [2006] recommends that prior expert knowledge is used to determine stable interactions. Therefore useful future work could be conducted to incorporate expert credit advice into the model build stage prior to automated modelling.

Finally, we ran a simple experiment using MV models for stress testing with hypothetical changes in interest rates. The results are plausible in the sense that the forecast mean LGDs were in the range we would expect given historic data. Nevertheless further work is needed to use models with MVs for stress testing, in terms of scenario generation and validation of stress test results.

Acknowledgements

We would like to thank our commercial partners for their assistance and comments in preparing this paper. Research was funded by UK EPSRC grant number EP/D505380/1, working as part of the Quantitative Financial Risk Management Centre.

References

Akaike H (1973). Information theory and an extension of the maximum principle. *Proc. 2nd Int. Symp. Information Theory*, Akademia Kiado, Budapest, 267-281.

Altman EI, Resti A and Sironi A (2005). Loss given default: a review of the literature. *Recovery Risk* ed. Altman E, Resti A and Sironi A (Risk Books).

Bank for International Settlements BIS (2005). Stress testing at major financial institutions: survey results and practice. *Working report from Committee on the Global Financial System*.

Basel Committee on Banking Supervision (BCBS 2006). *Basel II: International Convergence of Capital Measurement and Capital Standards*, Basel.

Bellotti T and Crook J (2007). Modelling and predicting loss given default for credit cards. *Quantitative Financial Risk Management Centre working paper.*

Brambor T, Clark WR and Golder M (2005). Understanding interaction models: improving empirical analyses. *Political Analysis* **14**: 63-82.

De Laurentis G and Riani M (2005). Estimating LGD in the Leasing Industry: Empirical Evidence from a Multivariate Model. In Altman E, Resti A and Sirona A (eds) *Recovery Risk*, Risk Books, London.

Dermine D and Neto de Carvalho C (2005). Bank loan losses-given-default: a case study. *Journal of Banking and Finance* vol.30, issue 4, 1219-1243

Figlewski S, Frydman H and Liang W (2007). Modeling the Effect of Macroeconomic Factors on Corporate Default and Credit Rating Transitions. *NYU Stern Finance working paper*. November 2007.

Financial Services Authority FSA (2005). Stress Testing. *Discussion paper* 05/2 FSA(UK).

Gayler R (2006). Comment on “Classifier technology and the illusion of progress – Credit scoring” by Hand D, *Statistical Science* vol 21, no 1, 19-23.

Grippa PS, Iannotti F and Leandri F (2005) Recovery rates in the banking industry: stylised facts emerging from the Italian experience. In Altman E, Resti A and Sirona A (eds) *Recovery Risk*, Risk Books, London.

Gupton GM and Stein RM (2005). *Losscalc V2: dynamic prediction of LGD, modelling methodology*. Moody's KMV Company.

Hand D (2006). Classifier technology and the illusion of progress. *Statistical Science* vol 21, no 1, 1-14.

Kennedy P (2003). *A Guide to Econometrics* (5th edition). Blackwell.
Lam J-P and Veall MR (2002). Bootstrap prediction intervals for single period regression forecasts. *International Journal of Forecasting* 18 pp 125-130.

Miller AJ (1990). *Subset Selection in Regression*. Chapman & Hall NY.

Querci F (2005). Loss Given Default on a medium-sized Italian bank's loans: an empirical exercise. *European Financial Management Association*
2005 (Milan, Italy).

Rosche D and Scheule H (2006) A multifactor approach for systematic default and recovery risk. In Engelmann B and Rauhmeier R (eds) *The Basel II Risk Parameters*. Springer, Berlin.

Saurina J and Trucharte C (2007). *An assessment of Basel II procyclicality in mortgage portfolios*. Banco de España.

Schuermann T (2005). What do we know about Loss Given Default? In Altman E, Resti A and Sirona A (eds) *Recovery Risk*, Risk Books, London.

Thomas LC, Edelman DB and Crook JN (2002). *Credit Scoring and its Applications*. SIAM Monographs on Mathematical Modeling and Computation. SIAM: Philadelphia, USA.