THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Intensity models and transition probabilities for credit card loan delinquencies

Citation for published version:

Leow, M & Crook, J 2014, 'Intensity models and transition probabilities for credit card loan delinquencies', *European Journal of Operational Research*, vol. 236, no. 2, pp. 685-694.
<https://doi.org/10.1016/j.ejor.2013.12.026>

Digital Object Identifier (DOI):

[10.1016/j.ejor.2013.12.026](https://doi.org/10.1016/j.ejor.2013.12.026)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

European Journal of Operational Research

Publisher Rights Statement:

© Leow, M., & Crook, J. (2014). Intensity models and transition probabilities for credit card loan delinquencies. *European Journal of Operational Research*. 10.1016/j.ejor.2013.12.026

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Intensity models and transition probabilities for credit card loan delinquencies

Abstract

We estimate the probability of delinquency and default for a sample of credit card loans using intensity models, via semi-parametric multiplicative hazard models with time-varying covariates. It is the first time these models, previously applied for the estimation of rating transitions, are used on retail loans. Four states are defined in this non-homogenous Markov chain: up-to-date, one month in arrears, two months in arrears, and default; where transitions between states are affected by individual characteristics of the debtor at application and their repayment behaviour since. These intensity estimations allow for insights into the factors that affect movements towards (and recovery from) delinquency, and into default (or not). Results indicate that different types of debtors behave differently while in different states. The probabilities estimated for each type of transition are then used to make out-of-sample predictions over a specified period of time.

Keywords: *risk analysis, probability of default, intensity modelling, time-varying covariates, state space modelling, retail loans*

1. Introduction

Risk models for retail portfolios of financial institutions, as well as within the academic literature, have not been developed as extensively as they have been in the corporate sector, mainly due to the availability and inaccessibility of the necessary data. However, with the financial crisis in 2008, awareness and the importance of credit risk management have increased, and new insights were gained, especially in terms of how correlated loans losses, debtor behaviour and the economic climate can be. In that, there has been work in the corporate sector, estimating Probability of Default (PD) and Loss Given Default (LGD) models with the inclusion of macroeconomic variables (for example, see Frye (2000a, 2000b) for PD; Gupton and Stein (2002, 2005) for LGD), but only recently has this been undertaken for retail loan credit models (for example, see Bellotti and Crook (2010), Pennington-Cross (2003) for PD; Bellotti and Crook (2012), Leow et al. (2011) for LGD).

Using a large dataset of credit card loan accounts provided by a major UK bank, we develop intensity models to predict delinquency and default. Our work differs from existing work in a number of ways. The majority of retail loans PD models currently in the literature are of static regression models (see Crook and Bellotti (2010), Leow and Mues (2012)), where models predicting default are developed using loan application characteristics and are valid only within a specified outcome period, e.g. within 12 months of opening. Such models are also unable to handle accounts that are active but have not (yet) experienced any event (known as censoring) or closed, so such accounts are usually deleted from the dataset used to develop such models. Furthermore, these models are only able to account for time-varying covariates at any single snapshot in time yet these indicators essentially change over time, so are unable to adequately incorporate the effect of macroeconomic variables. Subsequent work has been based on the use of survival models (see Banasik et al. (1999), Stepanova and Thomas (2002), Bellotti and Crook (2010)) for default risk, which will allow for a more dynamic prediction of events. Such models will predict not just the probability of whether an event will occur (and not limited to a pre-defined outcome period), but also the (conditional) probabilities of that event occurring over time. Although survival models can account for different types of events (via competing risks), they are based on the assumption that the risk of each event

occurring is independent right up to when any event occurs, which might not necessarily hold true. For example, in the months leading up to the default event, there would be certain behavioural indicators which take on values to indicate two or more missed monthly repayments, and this should mean that the risk of default increases more than the risk of early prepayment. There have also been papers suggesting the use of Markov chains (see Ho et al. (2004), Malik and Thomas (2012)), and although these have been useful in trying to quantify the behaviour of consumers, they have the complication of having to assume stationarity and finding the appropriate first, second or third order chain. We propose the use of intensity models to predict for delinquency and default, which have been previously applied to the estimation of rating transitions of corporate loans (see Jarrow et al. (1997), Duffie et al. (2007), Lando and Skodeberg (2002)), but have not been used on retail loans yet. Also, other approaches for estimating transition probabilities have been applied to corporate loans, for example the standard unobserved latent factor model and Bayesian methods (see for example Stefanescu (2009) and Kadam and Lenk (2008) but in these papers only aggregated data was used.

In this work, we do not just focus on the prediction of default. Instead, using both application and behavioural variables, we model time to delinquency, and then to default, based on how debtors have behaved throughout their loan period as well as how they might have handled previous experiences of periods in arrears. As such, we use an alternative definition of default here: three months of missed payments, but not necessarily consecutive; instead of the conventional one of defining default to occur when the debtor has missed three consecutive months (or 90 days' worth) of payments (The Financial Services Authority (2009), BIPRU 4.3.56 and 4.6.20). This then allows us to define four states chronicling the progression from up-to-date to default (to be defined in the following section), as well as when accounts move towards (or away) from default. Credit card accounts are tracked over a period of time where transitions between the various states could be affected by the individual characteristics of the debtor at time of application and how the debtor has managed their finances since gaining the credit account. Other external factors, like macroeconomic variables, could also be included but are not considered in this work. Each possible transition in this intensity model is modelled separately via a semi-parametric multiplicative hazard model with time-varying covariates (see Andersen et

al. (1993; 1991)), which are then calibrated to get the probabilities of moving (or staying) between states. Although this methodology has been detailed in a number of academic papers including Andersen et al. (1991), Jarrow et al. (1997), Lando and Skodeberg (2002) and Berd (2005) among others, they focus mainly on the estimation of parameter estimates, and where predictions were done, they were only in the time-homogenous case. The models we advance allow one to estimate a complete matrix of transition probabilities between any two repayment states between any two duration time periods for each case (in our application each account). By applying cut-offs they allow the prediction of the numbers of cases (in a sample) that would be expected to transit from one repayment state in one time period to any other state in any other time period the analyst may choose.

The models have several important advantages over cross-section regression models and over survival models. Compared with the former, the user can gain predictions of the probability of moving between states in any future time period, not just the time period chosen at the time of model estimation. Also, time varying covariates can be incorporated. Compared with simple survival models, intensity models give many additional types of predictions, for example predictions of entire transition probability matrices in any future time period for each borrower, rather than merely the hazard probability of default occurring in any specified duration time.

Intensity models have several potential uses by practitioners in financial institutions. First, since intensity models allow the predictions of transition probability matrices in any future time period, they would be crucially useful to the computation of a financial institution's economic capital in any future time period. Second, by developing a model that can predict the different states of delinquency, not only are we able to get predictions for default over time, we are also able to get more intricate predictions for *each state of delinquency* leading up to default. This would enable a lender to attain insights into the factors that affect movements towards, and recovery from, delinquencies, as well as factors contributing towards a move from delinquency into default. So a lender could identify those types of borrowers that are likely to progress to default, those likely to progress to merely two or just one payment in arrears and also those likely to *recover* from being in arrears. The models also allow the identification of *when* each borrower who is in arrears is likely

to recover. Third, although we do not go into detail here, this work could also be used by a lender to predict default risk in low- or zero-default portfolios. The analysis conducted on a low-default portfolio could underestimate default risk, but this might be mitigated by taking into account, by using the models we advance, the episodes where accounts go into arrears but not default.

The rest of this paper is structured as follows: The data and notation are described in section 2. Section 3 describes the methodology, and Sections 4 and 5 detail results and predictions respectively. Section 6 concludes.

2. Data and definitions

Data was supplied by a major UK bank and were active credit card accounts from all parts of the UK. This large dataset of more than 49,000 unique accounts is a random sample of credit cards that were issued from January 2002 up to June 2005, as well as their monthly histories since the account was opened, up to June 2006 or the time at which the credit card account was closed, whichever is earlier. Accounts that were still active in June 2006 are said to be censored at that time. As part of pre-processing, accounts that do not have consecutive monthly observations were removed, as well as accounts where their history did not start from the time the credit card was issued. Account and debtor characteristics available are common application variables, including type of employment, length of time the debtor had been with the bank, time at address and age; and behavioural variables available on a monthly basis, including spending and repayment amounts, credit limit and outstanding balance.

From accounts that were active during the period of May 2005 to June 2006, a random sample of unique accounts (20%), as well as all their respective monthly observations, were selected and kept separate as the validation sample. All other accounts were included in the training sample. Thus, the training and validation sets are kept completely separate.

2.1. Minimum repayment amount

A minimum repayment amount is required for the assignment of states but this information was not directly available from the provider of the data. We define τ as duration time since an account was opened and we define the minimum repayment amount in duration month τ , M_τ , to be the higher of 1% of the outstanding balance in month $\tau - 1$ or £5. It is possible for the minimum repayment amount to be £0, if there is zero outstanding balance on the account. Also, if the account is in credit, the minimum repayment amount required is also defined to be £0.

2.2. Definition of states

Four states are defined: state 0 means that the account is up to date; states 1 to 3 mean that the account is in one, two and three months in arrears, respectively. Note that these months in arrears do not necessarily have to be consecutive. State 3 is also known as the default state, so any account that enters state 3 is said to have entered default. For the purpose of this work, the default state is defined to be an absorbing state (i.e. accounts that enter the default state will not leave). States and movements between states are assigned as follows. Let M_τ denote the minimum repayment required from the statement from month τ , and let P_τ denote the repayment amount in month τ .

- All accounts start in state 0, where the account is said to have an up to date repayment schedule.
- At any time during the observation period, if repayment amount, P_τ , is less than the minimum, M_τ , required, the debtor shall advance into the next immediate state (e.g. if the account is in state 0 in month $\tau - 1$, and failed to meet the minimum repayment amount in month τ , it will be said to have moved to state 1 in month τ).
- A debtor who has missed a repayment before (i.e. either already in state 1 or 2) but manages to make some repayment amount, P_τ , in the following month(s) shall remain in that state if the repayment amount meets the

minimum repayment amount¹, i.e. $M_\tau < P_\tau < M_\tau + M_{\tau-1}$; or be moved to one lower state if the repayment amount meets the sum total of the minimum amount of the current and previous month but not the previous month's outstanding balance, $B_{\tau-1}$, i.e. $M_\tau + M_{\tau-1} < P_\tau < B_{\tau-1}$; or be moved to state 0 if the repayment amount is larger than the previous month's outstanding balance, i.e. $P_\tau > B_{\tau-1}$.

Figure 1: Observed transitions of three randomly selected accounts, A, B and C, from bottom to top.

As such, it is not possible for any account to advance more than one state at any one time interval, but it is possible for accounts to drop more than one state (e.g. a debtor who is in state 2 and manages to repay his loan fully will be said to move from state 2 to state 0 in that month). Figure 1 displays the observed transitions of three example accounts, over the duration time of each loan: account A (in the bottom

¹ It is possible that the payment the debtor missed (and hence moved states) is not the previous immediate month, so it would not be fair to look at the previous month's repayment amount to decide if the debtor will move to a lower state. For example, let's say the minimum repayments (and actual payment amounts and state) of months 4 and 5 are £56 (£0, move from state 0 to state 1) and £84 (£84, remain in state 1) respectively. In month 6, the minimum payment is £62 and the debtor makes a payment amount of £120. Although this amount should be enough to cover the minimum repayment of the missed payment in month 4 (£62+£56 < £120) and allow the debtor to move out of state 1 back to state 0, under our definitions, the debtor would have to repay at least the sum total of this and the previous month's minimum payment (i.e. £62+£84=£146) before he is said to move out of state 1 into state 0.

panel) went into state 1 twice before quickly going into default by month 16; accounts B and C (in the middle and top panels, respectively) did go into arrears but recovered, and both were up to date at the end of the sample period.

2.3. Distribution of default

The calculation of minimum repayment amount, the threshold for transition between states, and hence the transition to default, are defined in this work independent of the data provider. Also, the definition of default adopted here is such that an account is said to go into default once it goes three (not necessarily consecutive) months in arrears, so is not the conventional definition of default. The default rate at each time τ is then calculated as the number of accounts that go into default, at time τ as a proportion of the total number of active accounts, at time τ . Figure 2 gives an empirical distribution of this rate as observed in the training set, and we see that it is highest nearer the start of the loan and decreases with the age of the loan.

Figure 2: Distribution of defaults over time, based on training set.

3. Methodology

Before going into the details of the methodology, we note the issue of continuous and discrete time. One advantage of intensity models is the fact that they can be estimated in continuous time, however, given that the data we have are in monthly observations, and that we have a large dataset, we decide to approximate this to discrete time. The adjustment for continuous to discrete time is based on that described in Chen et al. (2005), but not always necessary if the dataset is large enough such that at least one event occurs at any given time or if the number of time periods under consideration is large.

The final output of the intensity model is the matrix of transition probabilities, $\mathbf{P}_i(s, t)$, where s and t are specific realisations of τ . The matrix of transition probabilities will, given an account's covariates, give the probabilities of transitions between each pair of states over a specified time period, s to t . In order to get this transition matrix, two other components have to be estimated: the transition intensity, which reflect the rate of change in the number of accounts between each pair of states at each time and are therefore both state and time dependent; and the time-dependent generator matrix, which is estimated via the cumulation of transition intensities up to time t . We note that in the case where there is time and account homogeneity, the estimation of the transition intensity, generator matrix and transition matrix can be estimated via the matrix exponential. However, this is not the case here so this model has to be estimated using a different method. These are further detailed below.

Following Andersen et al. (1991) and Lando and Skodeberg (2002), we let any individual i , $i = 1, \dots, n$, moving from state h to state j , $h, j = 0 \dots 3; h \neq j$, be represented by a vector of m covariates, $\mathbf{Z}_i(\tau)$, consisting of time-independent covariates (application variables) and time-dependent covariates (behavioural variables)². The transition intensity between states, α_{hji} , is then defined as in Equation 1.

² Note that $\mathbf{Z}_i(\tau)$ could be more specifically written as $\mathbf{Z}_{hji}(\tau)$, which would be specific to each particular h to j transition, but is not necessary in this work because the time-dependent variables are not state-specific.

$$\alpha_{hji}(\tau) = Y_{hi}(\tau) \alpha_{hj0}(\tau) \exp\{\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(\tau)\} \quad (1)$$

where $Y_{hi}(\tau)$ is an indicator for whether individual i was in state h at time τ , α_{hj0} is the baseline transition intensity for state h to state j and $\boldsymbol{\beta}_{hj}$ is a vector of unknown regression coefficients for the m covariates.

We estimate $\boldsymbol{\beta}_{hj}$ by $\hat{\boldsymbol{\beta}}_{hj}$, by maximising the generalised Cox partial likelihood (see Cox (1972) and Andersen et al. (1993), Section VII), given in Equation 2.

$$L(\boldsymbol{\beta}) = \prod_{\tau} \prod_{h,j,i} \left[\left(\frac{r(\boldsymbol{\beta}^T \mathbf{Z}_{hji}(\tau))}{S_{hj}^0(\boldsymbol{\beta}, \tau)} \right)^{\Delta N_{hji}(\tau)} \right] \quad (2)$$

where $r(\cdot)$ is some function to be defined, $S_{hj}^0(\boldsymbol{\beta}, \tau) = \sum_i^n [r(\boldsymbol{\beta}^T \mathbf{Z}_{hji}(\tau)) Y_{hi}(\tau)]$ and $\Delta N_{hji}(\tau) = 1\{N_{hji}(\tau) - N_{hji}(\tau - 1) = 1\}$, $N_{hji}(\tau)$ = number of observed transitions from state h to j by individual i over time τ .

We define function $r(\cdot)$ to be the exponential function, following Cox (1972) and Lando and Skodeberg (2002), and taking the logarithm of it, get Equation 3.

$$\begin{aligned} \log L(\boldsymbol{\beta}_{hj}) &= \log \left\{ \prod_{\tau} \prod_{h,j,i} \left[\left(\frac{\exp(\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(\tau))}{S_{hj}^0(\boldsymbol{\beta}_{hj}, \tau)} \right)^{\Delta N_{hji}(\tau)} \right] \right\} \\ &= \sum_{\tau} \sum_{h,j,i} \left\{ \Delta N_{hji}(\tau) [\log(\exp(\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(\tau))) - \log(S_{hj}^0(\boldsymbol{\beta}_{hj}, \tau))] \right\} \\ &= \sum_{\tau} \sum_{h,j,i} \left\{ \Delta N_{hji}(\tau) [\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(\tau) - \log(S_{hj}^0(\boldsymbol{\beta}_{hj}, \tau))] \right\} \end{aligned} \quad (3)$$

Where $S_{hj}^0(\boldsymbol{\beta}_{hj}, \tau) = \sum_i^n Y_{hi} \exp(\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(\tau))$.

The maximisation of the partial likelihood of Equation 3 will give values for $\hat{\boldsymbol{\beta}}_{hj}$ and estimates for $\boldsymbol{\beta}_{hj}$, which can then be used to calculate the baseline transition intensities, $\hat{A}_{hj0}(\tau) = \int_0^{\tau} \alpha_{hj0}(u) du$. These are estimated by the Nelson-Aalen type estimators and are given in Equation 4 (slightly altered because we have ties for

event times, see Keiding and Andersen (1989), Aalen et al. (2008), Chapter 4, and Borgan (1997)).

$$\begin{aligned}\hat{A}_{hj0}(\tau, \boldsymbol{\beta}_{hj}) &= \int_0^\tau \frac{J_h(u)}{S_{hj}^0(\hat{\boldsymbol{\beta}}_{hj}, u)} dN_{hj}(u) \\ &\cong \sum_0^\tau \frac{D_{hj}(u)}{S_{hj}^0(\hat{\boldsymbol{\beta}}_{hj}, u)}\end{aligned}\quad (4)$$

Where $J_h(u) = I\{Y_{h1} + \dots + Y_{hn} > 0\}$, $D_{hj}(u)$ = Number of type h to j transitions at time u ,

and $S_{hj}^0(\hat{\boldsymbol{\beta}}_{hj}, u) = \sum_i^n Y_{hi} \exp(\boldsymbol{\beta}_{hj}^T \mathbf{Z}_i(u))$, for $h, j = 0 \dots 3; h \neq j$.

In order to get the transitions probabilities, the generator matrix has to be defined. A generator matrix, $\mathbf{A}$, consists of the following elements.

$$\begin{aligned}\hat{A}_{hji}(\tau; \hat{\boldsymbol{\beta}}_{hj}, \mathbf{Z}_i(\tau)) &= \int_0^\tau \alpha_{hji}(u) du \\ &\cong \sum_0^\tau Y_{hi}(u) \alpha_{hj0}(u) \exp(\hat{\boldsymbol{\beta}}_{hj}^T \mathbf{Z}_i(u)) \quad h \neq j\end{aligned}\quad (5)$$

$$\hat{A}_{hhi}(\tau; \hat{\boldsymbol{\beta}}_{hj}, \mathbf{Z}_i(\tau)) = - \sum_{h \neq j} \hat{A}_{hji}(\tau) \quad h = 0 \dots 3 \quad (6)$$

Equation 5 cumulates the intensities of experiencing event hj up to each time τ , and the latter cumulates the intensities for any of the possible events.

The probabilities of transition from time s to t , for each individual i , $P(s, t, \mathbf{Z}_{hji}(\tau))$, given $\mathbf{Z}_{hji}(\tau)$, can be calculated by the product integral given in Equation 7. For simplicity, we rewrite $\mathbf{P}(s, t, \mathbf{Z}_{hji}(\tau))$ as $\mathbf{P}_i(s, t)$.

$$\begin{aligned}\hat{\mathbf{P}}_i(s, t) &= \prod_{(s, t]} (I + d\hat{\mathbf{A}}_i(u; \mathbf{Z}_i(u))) \\ &= \prod_{(s, t]} (I + \hat{\mathbf{A}}_i(u; \mathbf{Z}_i(u)) - \hat{\mathbf{A}}_i(u-1; \mathbf{Z}_i(u-1)))\end{aligned}\quad (7)$$

These transition matrices, which are specific to individuals, would give the probability of an account being in each state at time t , given the state it was in at time s . In the case of our dataset, where account observations are available monthly, we say that the probability of a transition at any time t is $\hat{\mathbf{P}}_i(t) = I + \hat{\mathbf{A}}_i(t; \mathbf{Z}_{hji}(t)) - \hat{\mathbf{A}}_i(t-1; \mathbf{Z}_{hji}(t-1))$.

4. Results

The intensities of six different transitions that are observed in the dataset are estimated here: from state 0 to 1, from state 1 to 2, from state 1 to 0, from state 2 to 3, from state 2 to 1 and from state 2 to 0. The final variables and their explanations included in the models are given in Table 1. Due to data confidentiality agreements, not all variables can be named.

4.1. Model parameter estimates

The covariates included in each transition model are kept the same for all the transitions, which would allow for some comparison of the effects of each covariate on the different transition types. For comparison purposes, only the parameter estimate signs are included and signs for all intensity models are listed in a single table, Table 1. The asterisks represent statistical insignificance at 0.05. From previous work (e.g. see Bellotti and Crook (2012)), we might have certain expectations of different types of debtors, for example, those employed would be expected to be most reliable, and younger debtors aged below 25 might be considered more risky than those in their 30s. However, it is with interest that we note that these prior expectations do not necessarily hold here, and that the effects of some covariates vary depending on the state the individual is in.

Most of the application variables have parameter estimate signs that behave consistently with previous literature (see Avery et al. (2004)), for example, a debtor that has been with the bank a longer time seem to have a lower chance of delinquency, and a higher chance of moving towards recovery; home owners are less likely to be delinquent and more likely to recovery should they go into delinquency. But there are interesting insights in two key application variables commonly used in scoring models: age and employment status. Age at application was divided into a number of groups with those 21 years old and below as the base category. Generally, it is observed that debtors in the older categories are less likely to move into states of delinquency and more likely to move out of delinquency, as compared to the base category. The exception is when debtors are already two months in arrears (state 2) here, younger debtors are less likely to actually default

(move from state 2 to 3), and more likely to move out of state 2. This might be because these debtors are able to eventually turn to their family for help. For employment status, those who are employed are fixed as the base category, which could be regarded as the preferred debtor due to the stability of income. Those that are self-employed or unemployed are more likely to go into states 1 and 2 (arrears), but not necessarily default. When looking at recovery from delinquency, the model implies that debtors that are employed are in fact less likely to recover, especially from state 2. We theorise that those that are self-employed or unemployed are probably used to struggling with their finances and therefore more adept at balancing their accounts, whereas those that are employed and who go into arrears are likely to have encountered an unforeseen situation and thus likely to slide into default quickly.

The time-dependent behavioural variables are lagged, so are known at time of prediction. They behave intuitively, for example, the higher the proportion of his/her available credit drawn, which could be an indicator of poor finances, the more likely the debtor will go into delinquency. Of interest is a variable representing the rate of transitions (RJT3), which could be an indication of the credibility of the debtor. It is observed that the higher the rate of jumps, the more likely the debtor will go into delinquency (i.e. states 1 or 2) but not default (i.e. state 3). This again implies that debtors who frequently go into delinquency (those that are used to struggling with their finances) seem to be able to balance their finances better. Also, the higher the rate of transitions, the less likely debtors are to make full recoveries (i.e. from state 1 to 0, or state 2 to 0) but more likely to make some payment that will move them to a lower state of delinquency (i.e. state 2 to state 1).

To check for model fit, martingale residuals could be computed, but with the inclusion of time-dependent covariates, the underlying assumptions of the residual calculations no longer hold. Therefore, we validate the model by comparing the number of predicted and observed transitions, as detailed in Section 5.2.

Table 1. Parameter estimate signs for the intensity models of different types of transitions. The asterisk represents a statistical insignificance at 0.05.

Code	Explanation	Towards delinquency			Towards recovery		
		0 to 1	1 to 2	2 to 3	1 to 0	2 to 1	2 to 0
NOCA	Number of cards at application	-	+	+	-	-	+
LAND	Indicator for presence of landline	-	-	-	+	+	+
TADD	Time at address, in years	+	-	-	+	+	+
TWBA	Time with Bank, in months	-	-	-	+	+	+
TWBM	Indicator for missing Time with Bank	-	-	-	+	+	+
INCL	Income, Ln	-	+	+	-	-	+
INCM	Indicator for missing income	-	+	+	-	-	+
A1	Variable A Group 1	.	.	.	.	.	.
A2	Variable A Group 2	+	+	+	-	-	-
A3	Variable A Group 3	+	+	+	-	-	-
A4	Variable A Group 4	+	+	+	-	-	-
A5	Variable A Group 5	+	+	+	-	-	-
AGE1	Age at application, group 1	.	.	.	.	.	.
AGE2	Age at application, group 2	-	-	+	+	-	-
AGE3	Age at application, group 3	-	-	+	+	-	-
AGE4	Age at application, group 4	-	-	+	+	-	-
AGE5	Age at application, group 5	-	-	+	+	-	-
AGE6	Age at application, group 6	-	-	+	+	-	-
AGE7	Age at application, group 7	-	-	+	+	-	-
AGE8	Age at application, group 8	-	-	+	+	-	-
AGE9	Age at application, group 9	-	-	+	+	-	-
AGE10	Age at application, group 10	-	-	+	+	-	+
EEMP	Employment Group: Employed	.	.	.	.	.	.
ESEL	Employment Group: Self-employed	+	+	+	-	+	-
ENOT	Employment Group: Not employed	-	+	+	-	-	+
EUNE	Employment Group: Unemployed	+	+	-	+	+	+
CLI3	Credit limit, ln, lagged 3 months	+	-	-	+	+	+
PAY3	Payment amount, ln, lagged 3 months	+	+	-	+	+	+
PDR3	Proportion of credit drawn, lagged 3 months	+	+	+	-	+	-
RJT3	Rate of total jumped, lagged 3 months	+	+	-	-	+	-
RSD3	Indicator for improvement in state from 3 months previous, lagged 3 months	+	+	-	-	+	+

4.2. Baseline intensity

The baseline intensity was calculated via Equation 4 for all the transitions. In order to make comparisons, we overlay the graphs for the transitions in which the underlying risk sets are common, for example, an account currently in state 1 could move to state 0 or move to state 2, so transitions from 1 to 2 and 1 to 0 would have a common risk set. Figures 3 to 5 represent the baseline intensities for transitions from states 0, 1 and 2 respectively. All three figures have graphs that start from month 4 because a three-month lag is applied for the behavioural variables.

Figure 3. Baseline intensity for transition from state 0 to state 1.

In Figure 3, we see the baseline intensity for transition from state 0 to 1 is highest near the beginning of the loan, which could represent accounts or debtors who were struggling with debt and go into delinquency soon after the credit card account is approved. The intensity then tapers off and remains quite stable throughout the

period of the loan, which would represent a small percentage of debtors who occasionally miss a payment unintentionally.

Figure 4. Baseline intensity for transitions from state 1. The stars represent the baseline transition intensity from state 1 to state 2; the squares represent the baseline transition intensity from state 1 to state 0.

Figure 4 displays the baseline intensity of two transitions, both transitions from state 1, to either state 0 or state 2. We observe that the transition intensity of state 1 to 0 is higher than the transition intensity of state 1 to 2; in other words, for accounts that are in state 1, there is a higher probability of moving towards recovery (state 0) than of moving towards further delinquency (state 2). Also, the transition intensity to state 2 is higher at the beginning of the loan, which implies that some debtors struggle from the very beginning and so go from state 1 to state 2 in a short period. With the exception of the couple of months near the beginning of the loan, transition intensities from state 1 are relatively flat throughout the duration of the loan.

Figure 5. Baseline intensity for transitions from state 2. The stars represent the baseline intensity from state 2 to state 3; the squares represent the baseline intensity from state 2 to state 1; the triangles represent the baseline intensity from state 2 to state 0.

Figure 5 displays the baseline intensities of transitions from state 2. The transition intensity from state 2 to state 3 is highest, followed by that to state 1, and transition intensity from state 2 to state 0 is lowest. This would mean that, contrary to what was observed in the previous graph on transitions from state 1, debtors that are in state 2 are more likely to go into default (state 3) than to make some recovery (state 1), and have an even lower chance of making full recovery (state 0).

5. Predictions

This model was developed such that, given the covariates of an individual account at any particular time, a matrix of transition probabilities of moving between states can be computed for any specified time period up to the length of lag of the behavioural

variables³. Two kinds of predictions are made here in order to first, get an insight of how the model would predict for accounts with different characteristics, and second, to validate the model using predictions.

5.1. Insights based on employment type

An example account for each employment type is created. The application variables are the mean (median or mode where appropriate) values of the accounts in each employment category from the training set. The behavioural variables at each time point are the mean (or mode) values of all active accounts at that time point. We limit the prediction time period to a 6 month period, 6 months after the account was opened⁴, i.e. the probabilities of transiting or staying in certain states at the end of month 12, given that the account was in a certain state in month 6. These matrices are given below.

$$\begin{aligned}
 \text{Employed (6, 12)} &= \begin{pmatrix} 0.8518 & 0.1228 & 0.0154 & 0.0098 \\ 0.7835 & 0.1294 & 0.0267 & 0.0602 \\ 0.3917 & 0.0843 & 0.0398 & 0.4841 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\
 \text{Self employed (6, 12)} &= \begin{pmatrix} 0.8707 & 0.1118 & 0.0119 & 0.0054 \\ 0.8261 & 0.1193 & 0.0201 & 0.0343 \\ 0.5426 & 0.0990 & 0.0345 & 0.3237 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\
 \text{Unemployed (6, 12)} &= \begin{pmatrix} 0.8583 & 0.1218 & 0.0125 & 0.0072 \\ 0.8157 & 0.1247 & 0.0197 & 0.0397 \\ 0.4611 & 0.0907 & 0.0370 & 0.4110 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\
 \text{Not employed (6, 12)} &= \begin{pmatrix} 0.8364 & 0.1352 & 0.0175 & 0.0108 \\ 0.7803 & 0.1375 & 0.0264 & 0.0555 \\ 0.4207 & 0.0958 & 0.0395 & 0.4437 \\ 0 & 0 & 0 & 1 \end{pmatrix}
 \end{aligned}$$

³ Predictions further into the future can be made by first forecasting the values of behavioural variables.

⁴ Prediction is done from six months because the observations from the first three months are not reliable across all observations, and because there is a 3-month lag for the behavioural variables.

A few of observations are made. If the account is originally in state 0, there is very high chance (above 80%) that it will still be in state 0 in month 12. Note that this does not necessarily mean that it is in state 0 for the entire 6 month period, it could have moved out of state 0 and then back in again. The chance of it being in state 1, 2 or 3 at month 12 decreases appropriately and this applies for all employment types. If the account is originally in state 1, there is a high chance (around 80% but lower than if it were in state 0, ranging from 78% to 82%) that it will have gone to state 0 by month 12. The next highest probability would be to stay in state 1. However, the probability of being in state 3 at month 12 is higher than the probability of being in state 2. This seems to suggest that for an account in delinquency, there is a high chance it will recover (back to 0), there is also a chance that it will just remain in state 1 but in the case where it is going to go bad, it is more likely to move to state 3 (default), then to be hanging on in state 2. If the account is originally in state 2, there is very low chance it will be in state 1 or 2 at month 12, and roughly the same chance of the account being in state 0 (full recovery) and in state 3 (default). The probabilities are different for the different employment types. For employment, there is a higher chance of being in default (48%) than recovery (39%); for self employed, a higher chance to be in recovery (54%) than in default (32%) and the difference in the two probabilities is much bigger than that seen for the employed; for unemployed, roughly the same at 43% either way; for not employed, higher chance of recovery (46%) than default (41%), but the difference in probabilities is not as great as for the self-employed. This seems to suggest that the self-employed, unemployed and not employed are used to struggling with balancing their finances and have a good chance of recovery, whereas for the employed people, they are likely to go into disaster mode.

5.2. Validation of test set

There are over 40,000 unique accounts in the test set that have at least six months' observations, for which we have application variables, monthly behavioural variables up to the time the account is closed or June 2006, whichever is earlier, including the state in which the account is in at each time point.

5.2.1. Test framework

The six transition models developed are applied to all accounts in the test set, from which the individual-specific transition matrices, $\mathbf{P}_i(s, t)$, are computed for time period s to t , the format of which is given in Equation 8. Although it is possible to get probabilities of transitions for any time point, it is necessary to define a period of time s to t in our work, so as to enable the comparison of the predicted state at time t with that of the actual state each account is in. In this work, we find the transition matrix for time 6 to 12, i.e. given states of accounts at time 6, to predict its state at time 12.

$$\mathbf{P}_i(s, t) = \begin{pmatrix} p_{00} & p_{01} & p_{02} & p_{03} \\ p_{10} & p_{11} & p_{12} & p_{13} \\ p_{20} & p_{21} & p_{22} & p_{23} \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (8)$$

In order to predict whether a transition has taken place, we compare predicted transition probabilities to cut-off values, which will be different for the different transitions. The cut-off value for each type of transition, c_{hj} , is computed such that the proportion of accounts predicted to undergo transition hj in the test set is equal to the proportion of accounts that undergo transition hj in the training set. However, because some accounts start less than 12 months before the end of our sample period, their states are unknown at time 12, and although it is possible to identify and remove these accounts from the test and training sets, these accounts are not usually known at the time of prediction. Note that predictions can still be made for these accounts, but how these accounts are handled will affect the proportion of observed transitions and hence the cut-off values. One way to handle this would be to remove accounts that have states unknown at time 12, so that these would not impact cut-off values and performance measures, known as scenario A. As a test of robustness, we consider two further scenarios: (B) let the state of the accounts at time 12 be the initial state of the account, i.e. at time 6; and (C) let the state of the accounts at time 12 be the state each account was last observed in. Each scenario will produce a different number of transitions and different cut-off values. Given the observed state of each account at time 6 and the probabilities from the transition matrix, the predicted state of account i at time 12, is given by equation 9.

The order of which the states are compared affects the predictions because the states are competing states, i.e. an account predicted to go into the first state compared, would not be considered for transition into the rest of the states. For example, by comparing the probabilities and cut-off values of transition hj_1 before transition hj_2 , we place an additional (and random) chance of transition to state j_1 over state j_2 because should a transition to state j_1 happen, the account is no longer at risk for a transition to state j_2 . By assigning one state before another, we implicitly place a higher chance of transition on the first state, which is debatable as an equally compelling reason could be found for a different order. In this work, we rank states based on the training set, in ascending order of number of transitions, and compare probabilities in that order⁵. Note that this ordering could be different for different states as well as for the different scenarios.

$$\text{predict}(\text{state}_i(t) | \text{state}_i(s) = h) = \begin{cases} j_1 & \text{if } p_{hj_1} > c_{hj_1} \\ j_2 & \text{if } p_{hj_1} \leq c_{hj_1} \text{ and } p_{hj_2} > c_{hj_2} \\ j_3 & \text{if } p_{hj_1} \leq c_{hj_1} \text{ and } p_{hj_2} \leq c_{hj_2} \text{ and } p_{hj_3} > c_{hj_3} \\ j_4 & \text{otherwise} \end{cases} \quad (9)$$

Where states j_1 to j_4 represent states 0 to 3.

5.2.2. Performance measures

With more than two states defined in this work, the usual predictive analytics performance measures cannot be applied here. For each scenario, the predicted states are then compared against the observed states at time t , and three different performance measures are computed. The first is the cohort level accuracy table, which gives the overall number of accounts predicted to be in each state at time 12 (columns), given the number of accounts in each state at time 6 (rows). Due to confidentiality agreements with the data provider, the actual numbers cannot be reported, and each cell is calculated to be the ratio $\frac{\text{number of predicted } hj \text{ transitions}}{\text{number of observed } hj \text{ transitions}}$.

As such, a value greater than 1 would mean that the number of predicted hj transitions are greater than the number of observed hj transitions, i.e. the number of accounts predicted to make this particular transition has been over-estimated.

⁵ An alternative ranking of states was considered with similar results and performance measures.

Similarly, a value less than 1 would imply that the number of accounts predicted to make this particular transition has been under-estimated. The ratio itself will also give an indication of how far off the predictions are, for example a value of 2 would that the number of predicted transitions is double what is observed. From Table 2, we see that overall predicted numbers are quite close to what was observed, and the model tends to over-estimate the number of accounts going into delinquencies; and in the case of defaults, these numbers are very much over-estimated. For example, from accounts that were in state 2 at time 6 (scenario A), the model predicts 11 times the number of accounts that will default (state 3), and only a third of the number of accounts that made a full recovery (state 0) by month 12. When accounts with states unknown at time 12 are included in the predictions (scenarios B and C), we see that while the ratios in the other cells remain similar to the ratios in scenario A, the ratios of defaults increases further, maintaining the conservative predictions.

Next is the account level confusion table, where the rows and columns of this table give the numbers of observed and predicted accounts for each state at time 12, respectively, and the proportion of states that are actually correctly predicted. As before, it is not possible to report actual numbers, and cells are divided by different sums in order to conceal the actual delinquency and default rates. The diagonal cells (in bold) of each panel in Table 3 give the sensitivity of each state, representing the proportion of accounts that have their states at time 12 correctly predicted. These ratios show that the model is able to distinguish between delinquent and non-delinquent accounts well, with an accuracy percentage of above 91% for state 0 predictions. However, it is less successful at differentiating between delinquent states, with accuracy percentages of 15% to 20% for state 3 predictions. The non-diagonal cells represent the proportion of accounts that are wrongly predicted, and calculated to be the ratio $\frac{\text{number of accounts in state } h \text{ but predicted to be in state } j}{\text{total number of accounts predicted to be in state } j}$. The cells above the diagonals represent accounts that have been predicted to be in a state worse than it actually is, while those below the diagonals represent accounts that have been predicted to be in a better state. The difference in the ratios in these two groups of cells are striking, which show that, where the model gets the predictions wrong, it is conservatively skewed, with the majority of wrong predictions

to be in a state higher than they actually are. This result holds across the different scenarios.

Table 2: Cohort level accuracy table.

		Predicted at time 12											
		(A) Unknowns left as is				(B) Unknowns as initial state				(C) Unknowns as last known state			
		State 0	State 1	State 2	State 3	State 0	State 1	State 2	State 3	State 0	State 1	State 2	State 3
Observed at time 6	State 0	1.004	0.696	2.132	12.829	0.998	0.751	2.229	13.610	1.004	0.713	2.178	13.268
	State 1	0.907	0.614	1.563	20.036	1.004	0.569	1.676	21.571	0.919	0.651	1.432	21.107
	State 2	0.365	0.327	0.647	11.600	0.397	0.347	0.404	12.800	0.377	0.351	0.500	12.933
	State 3	0	0	0	0	0	0	0	0	0	0	0	0

Table 3: Account level confusion table.

		Predicted at time 12											
		(A) Unknowns left as is				(B) Unknowns as initial state				(C) Unknowns as last known state			
		State 0	State 1	State 2	State 3	State 0	State 1	State 2	State 3	State 0	State 1	State 2	State 3
Observed at time 12	State 0	0.910	0.798	0.749	0.780	0.914	0.773	0.741	0.747	0.913	0.804	0.740	0.777
	State 1	0.078	0.117	0.173	0.173	0.077	0.138	0.175	0.196	0.076	0.117	0.180	0.171
	State 2	0.005	0.024	0.076	0.034	0.004	0.020	0.083	0.047	0.005	0.024	0.085	0.040
	State 3	0.001	0.004	0.038	0.202	0.001	0.005	0.039	0.155	0.001	0.004	0.034	0.190

Finally, we tabulate some overall statistics, one of which is the accuracy percentage, which will give the percentage of accounts with correctly predicted states at time 12. We also calculate the proportion of accounts that have been predicted to be in a state worse than it should be, i.e. the model is being conservative; and correspondingly, the proportion of accounts that have been predicted to be in a state better than it should be, i.e. the model is being too optimistic. These are given in Table 4, for all three scenarios.

Table 4: Overall statistics

Scenario	Accuracy	Too conservative	Too optimistic
(A) Unknowns left as is	83.25%	9.05%	7.70%
(B) Unknowns as initial state	83.77%	8.64%	7.59%
(C) Unknowns as last known state	83.81%	8.68%	7.51%

Overall, the results show that the model is able to predict quite accurately, achieving an accuracy percentage of around 83% (cf. Table 4), but this high accuracy percentage figure is probably due to the high accuracy for predictions in state 0, and closer inspection of predictions for accounts that were initially in states 1 and 2 at the beginning of the test periods are very poor. Although we do not have the relevant costs for each type of wrong prediction, the cost associated with a default wrongly predicted to be a recovery is expected to be greater than the cost associated with a recovery wrongly predicted to be a default. From Table 4, we see that the model does give conservative predictions, as a larger percentage of accounts are predicted to be in a worse state than they are. For example, from Table 3, panel A, 78% of accounts that were predicted to be in state 3 were actually in state 0, whereas 0.1% of accounts that were predicted to be in state 0 were actually in state 3. The cost of misclassifying each of the former accounts is greater than that of misclassifying each of the latter accounts. Thus the model's predictions over a 12 month horizon overestimated the costs of misclassifying accounts rather than underestimated them. Notice also that the model appears robust as it gives similar results across the different scenarios.

5. Conclusions

Based on a large dataset of credit card loans, we developed a set of semi-parametric multiplicative intensity models to predict delinquency. These models, which are based on survival models, are able to incorporate time-dependent variables, and are able to predict not just whether an event will occur, but also give probabilities of when it might occur, thus providing a more dynamic framework for prediction. It is the first time these models are being applied to retail loans, and we were able to achieve two main research outcomes in this work.

First, interesting insights into the factors that affect movements towards (and recovery from) delinquency were made. By keeping the covariates unchanged for the different types of transitions, we were able to compare the effects of each covariate on the different transitions. We found that most application variables affect risk of delinquency similarly to what was previously established (based on credit scoring or behavioural models), but also that some groups of people are better in keeping themselves only in delinquency without tipping over to default. In particular, the self-employed and unemployed are at higher risk of going into delinquency as compared to those employed; but once in delinquency, the employed seem less able to recover and avoid the state of default. This phenomenon was again seen in one of the time-dependent behavioural variables, where we observed that those who frequently go into delinquency are more likely to go into delinquency again but not default. We theorised that debtors who are self-employed or unemployed are better at balancing their accounts, and thus more able to stay out of default, whereas those who are employed enjoy a stable income and thus are less successful at keeping out of default when confronted with an unexpected break in their income. The baseline intensities computed were also intuitive for all transitions.

Secondly, the model was used for predictions in two parts, one to gain insights based on employment type, and the other to validate the model. In the first part, a typical account for each employment type was created and the matrix of transition probabilities was computed for time 6 to 12 months after the account was opened. We found that the model produced plausible and intuitive results, where in general, accounts that were non-delinquent at time 6 are very likely to remain non-delinquent at

the end of month 12, but did have a small chance of going into one of the different states of delinquency; and how delinquent the debtor was at time 6 would affect his chance of recovery, further delinquency or default. We observe that debtors of different employment types behaved differently when in different states. In particular, while in delinquency, debtors who are self-employed or unemployed seemed more adept at keeping themselves only in arrears without tipping over into default.

In the second part, we applied the intensity models to all observations in the test set at time 6 months into the start of the account and, using an algorithm and different cut-off values for different transitions, predicted for the state of each observation at time 12 months. Three different scenarios were considered for the handling of accounts that were censored before the end of the observation period. From these predictions, we found that the model made fairly accurate predictions on an overall level, but was not able to do as well on the account level. The total number of accounts predicted to be in each state at the end of the respective test periods were reasonably close to what was observed, but while the model was able to predict for transitions, it did not seem able to predict the correct accounts transit. Overall, we found that the model is adept at predicting between delinquent and non-delinquent accounts as a high percentage of accounts were correctly predicted to be in state 0 at time 12; but did not do as well when predicting between delinquent states, i.e. states 1, 2 and 3. When the model got it wrong, it erred on the conservative side, by predicting the account to be in a state lower than it should be, than liberal, which is expected to have a lower cost consequence.

We believe there is more research to be done in this area, especially since this is the first time intensity models are being applied to retail loans. One obvious addition to the model is time-dependent macroeconomic variables. While it is true that all debtors will be affected by changes in the economy, it is likely that different groups of debtors are affected to a variety of extents. An extension of this work would be to include some interaction terms between macroeconomic variables and some of the application covariates, as well as applying the model on more recent data, which could hopefully encompass the credit crisis of 2008. Another area of further work would be the way the probabilities of the model are translated into predicted events. Right now an algorithm involving cut-off values is used, and comparison of probabilities and cut-off

values is based on a particular specified order of states, and this implicitly increases the likelihood of transition into one state over another. Further research would be to experiment with different cut-off values, or possibly a different method of predicting for transition.

References

- Aalen, O. O., Borgan, Ø., & Gjessing, H. (2008). *Survival and Event History Analysis*: Springer.
- Andersen, P. K., Borgan, Ø., Gill, R. D., & Keiding, N. (1993). *Statistical Models Based on Counting Processes*: Springer-Verlag New York.
- Andersen, P. K., Hansen, L. S., & Keiding, N. (1991). Non-and Semi-Parametric Estimation of Transition Probabilities from Censored Observation of a Non-Homogeneous Markov Process. *Scandinavian Journal of Statistics*, 18, 153-167.
- Avery, R. B., Calem, P. S., & Canner, G. B. (2004). Consumer credit scoring: Do situational circumstances matter? *Journal of Banking & Finance*, 28, 835-856.
- Banasik, J., Crook, J. N., & Thomas, L. C. (1999). Not if but When will Borrowers Default. *The Journal of the Operational Research Society*, 50, 1185-1190.
- Bellotti, T., & Crook, J. N. (2010). Retail Credit Stress Testing using a Discrete Hazard Model with Macroeconomic Factors. *Credit Research Centre Working Paper* 10/2.
- Bellotti, T., & Crook, J. N. (2012). Loss given default models incorporating macroeconomic variables for credit cards. *International Journal of Forecasting*, 28, 171-182.
- Berd, A. M. (2005). Dynamic Estimation of Credit Rating Transition Probabilities.
- Borgan, Ø. (1997). Three Contributions to the Encyclopedia of Biostatistics: The Nelson-Aalen, Kaplan-Meier, and Aalen-Johansen Estimators.
- Chen, L.-S., Yen, M.-F., Wu, H.-M., Liao, C.-S., Liou, D.-M., Kuo, H.-S., & Chen, T. H.-H. (2005). Predictive survival model with time-dependent prognostic factors: development of computer-aided SAS Macro program. *Journal of Evaluation in Clinical Practice*, 11, 181-193.
- Cox, D. R. (1972). Regression Models and Life-Tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34, 187-220.

- Crook, J. N., & Bellotti, T. (2010). Time Varying and Dynamic Models for Default Risk in Consumer Loans. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 173, 283-305.
- Duffie, D., Saita, L., & Wang, K. (2007). Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics*, 83, 635-665.
- Financial Services Authority. (2009). Prudential Sourcebook for Banks, Building Societies and Investment Firms.
- Frye, J. (2000a). Collateral Damage. *RISK*, 13.
- Frye, J. (2000b). Depressing Recoveries. *RISK*, 13.
- Gupton, G. M., & Stein, R. M. (2002). LossCalc: Moody's Model for Predicting Loss Given Default (LGD). *Moody's KMV Company*.
- Gupton, G. M., & Stein, R. M. (2005). LossCalc v2: Dynamic Prediction of LGD. *Moody's KMV Company*.
- Ho, J., Thomas, L. C., Pomroy, T. A., & Scherer, W. T. (2004). Segmentation in Markov Chain Consumer Credit Behavioural Models. In L. C. Thomas, J. N. Crook & D. B. Edelman (Eds.), *Readings in Credit Scoring*. Oxford: Oxford University Press.
- Jarrow, R. A., Lando, D., & Turnbull, S. M. (1997). A Markov model for the term structure of credit risk spreads. *Review of Financial Studies*, 10, 481-523.
- Kadam, A., & Lenk, P. (2008). Bayesian inference for issuer heterogeneity in credit ratings migration. *Journal of Banking & Finance*, 32, 2267-2274.
- Keiding, N., & Andersen, P. K. (1989). Nonparametric Estimation of Transition Intensities and Transition Probabilities: A Case Study of a Two-State Markov Process. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 38, 319-329.
- Lando, D., & Skødeberg, T. M. (2002). Analyzing rating transitions and rating drift with continuous observations. *Journal of Banking & Finance*, 26, 423-444.
- Leow, M., & Mues, C. (2012). Predicting Loss Given Default (LGD) for Residential Mortgage Loans: A Two-stage Model and Empirical Evidence for UK Bank Data. *International Journal of Forecasting*, 28, 183-195.
- Leow, M., Mues, C., & Thomas, L. C. (2011). The Economy and Loss Given Default: Evidence from Two UK Retail Lending Datasets. *Working Paper*.
- Malik, M., & Thomas, L. C. (2012). Transition matrix models of consumer credit ratings. *International Journal of Forecasting*, 28, 261-272.

- Pennington-Cross, A. (2003). Subprime and Prime Mortgages: Loss Distributions. *OFHEO Working Paper 03-01*.
- Stefanescu, C., Tunaru, R., & Turnbull, S. (2009). The credit rating process and estimation of transition probabilities: A Bayesian approach. *Journal of Empirical Finance*, 16, 216-234.
- Stepanova, M., & Thomas, L. C. (2002). Survival Analysis Methods for Personal Loan Data. *Operations Research*, 50, 277-289.