European Journal of Operational Research 254 (2016) 977–984

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Modelling credit grade migration in large portfolios using cumulative t-link transition models☆

Jonathan J. Forster[a,*], Matteo Buzzacchi[b], Agus Sudjianto[c], Risa Nagao[d]

[a] Southampton Statistical Sciences Research Institute, University of Southampton, Highfield, Southampton SO17 1BJ, UK
[b] Barclays Capital, UK
[c] Wells Fargo, USA
[d] University of St Andrews, UK

ARTICLE INFO

Article history:
Received 10 January 2014
Accepted 11 March 2016
Available online 7 April 2016

Keywords:
Markov processes
Cumulative link
Heavy-tailed
Logistic
Probit

ABSTRACT

For a credit portfolio, we are often interested in modelling the migration of accounts between credit grades over time. For a large retail portfolio, data on credit grade migration may be available only in the form of a series of (typically monthly) population transition matrices representing the gross flow of accounts between each pair of credit grades in the given time period. The challenge is to model the transition process on the basis of these aggregate flow matrices. Each row of an observed transition matrix represents a sample from an ordinal probability distribution. Following Malik and Thomas (2012), Feng, Gourieroux, and Jasiak (2008) and McNeil and Wendin (2006), we assume a cumulative link model for these ordinal distributions. Common choices of link function are based on the normal (probit link) or logistic distributions, but the fit to observed data can be poor. In this paper, we investigate the fit of alternative link specifications based on the t-distribution. Such distributions arise naturally when modelling data which arise through aggregating an inhomogeneous sample of obligors, by combining a simple structural-type model for credit migration at the obligor level, with a suitable mixing distribution to model the variability between obligors.

1. Introduction

1.1. Background

Credit ratings can be an invaluable tool for describing and modelling the default risk for obligors in a particular loan pool. A credit rating system is an ordinal classification reflecting the probability of default of a given obligor, with the highest rating representing lowest probability of default, ranging down to a lowest rating typically representing an obligor already in default. For corporate assets, various ratings agencies (for example Moody's or Standard and Poor's) provide credit ratings. For retail obligors, a bank typically uses its internal credit scoring systems for rating purposes. Our work in this area is motivated by the need to model credit grade transitions in retail portfolios. Hence, we assume that grades are only observed at fixed discrete time intervals, and that detailed obligor-level covariate information is not available. However, we observe the same behaviour in retail portfolios and in pooled corporate agency ratings, and it is the latter which we use to illustrate our proposed modelling approach.

One possible method for forecasting the evolution of default risk in a portfolio is to forecast the process by which individual credit grades (including default) migrate over time. The natural description of the process of credit grade migration between two time points is the transition matrix $P(t, u)$, with elements $p_{ij}(t, u)$, representing the probability of transition from grade i at time point t to grade j at time point u, that is

$$p_{ij}(t, u) = \text{Prob}(\text{grade at time } u = j \,|\, \text{grade at time } t = i).$$

Here, we assume that i and j run from 1 (highest quality) to D with the final grade D representing default, and that, as described above, the grades are naturally ordinal with increasing grade number representing increasing closeness to default.

If $\pi(t) = \{\pi_1(t), \dots, \pi_D(t)\}$ is the row vector containing the proportions of obligors in each of the credit grades at time t, then a forecast of the corresponding proportions $\pi(u)$ at time $u > t$ is

☆ This article represents the views and analysis of the authors only and should not be taken to represent those of any of their current or former employers. Much of this work was carried out while the first three authors were colleagues, and the fourth author a project student, at Lloyds Banking Group. The authors are grateful for the helpful comments of two anonymous reviewers on an earlier version of this paper.

* Corresponding author. Tel.: +44 23 80595130.
E-mail address: J.J.Forster@soton.ac.uk (J.J. Forster).

http://dx.doi.org/10.1016/j.ejor.2016.03.017
0377-2217/

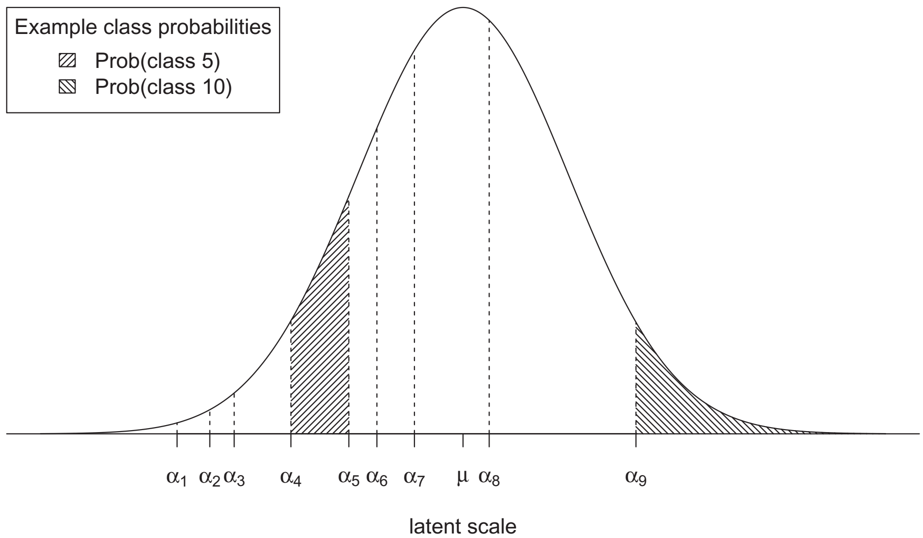

Fig. 1. Mapping between threshold and mean parameters and class probabilities for a cumulative link model.

given by

$$\pi(u) = \pi(t)P(t, u).$$

Hence, forecasting a future credit grade profile can be achieved by estimating the corresponding transition matrix between the present and the forecast horizon. Estimators are constructed by developing statistical models for transition matrices, and fitting them to observed data on historical transitions. In this paper, we focus on developing statistical models which fit historical transition data on portfolio credit grade distributions. Typically, these data are a series of portfolio transition matrices representing the gross flow of obligors between each pair of credit grades in the given time period. We do not consider here the question of modelling the complete transition process. That is, given an empirical matrix $X(t, u)$ of portfolio flow between times t and u, we consider the problem of estimating the corresponding transition matrix $P(t, u)$ which can be thought of as a smoothed (and normalised so that rows sum to one) version of $X(t, u)$. The matrices $X(t, u)$ and $P(t, u)$ are $D \times D$ matrices, with elements $x_{ij}(t, u)$ representing the observed number of transitions from state i to state j between times t and u, and $p_{ij}(t, u)$ the corresponding transition probability. As the default state D is an absorbing state, estimation of $P(t, u)$ simply requires estimation of the first $D - 1$ rows of $P(t, u)$.

We model the transition process over time as a series of one-period models, rather than a single all-encompassing model for the panel of empirical transition matrices. For a complete forecasting model, it is necessary to augment these one-period models for individual historical transitions, in order to predict future transition dynamics. This is discussed briefly in Section 5 and is the subject of ongoing research. Henceforth, for clarity, the dependence of P, (p_{ij}) and X (x_{ij}) on t and u, the beginning and end points of the period under consideration, is considered as implicit, and omitted from our notation.

1.2. Ordinal data models

Ordinal data can often be effectively modelled using a cumulative link model. A cumulative link model, for a collection of ordinal variables $\{Y_k\}$, can be written as

$$P(Y_k \le j) = g(\alpha_j - \mu_k) \tag{1}$$

for some strictly increasing function g which can be interpreted as the distribution function of a latent continuously distributed variable. Then, $-\infty = \alpha_0, \alpha_1, \ldots, \alpha_D = \infty$, can be thought of as an increasing sequence of thresholds, defining the mapping between the underlying latent scale and the ordinal classes. The μ_k parameters are observation-specific, but are typically modelled using a parsimonious regression function. Common choices of g are based on the standard normal (ordinal probit model) or logistic distributions (proportional odds model). A visual illustration of the mapping between the threshold parameters $\alpha_1, \ldots, \alpha_{D-1}$, the mean parameter μ and the class probabilities is given in Fig. 1 (for the case where $D = 10$).

For ordinal transition matrix modelling, each row represents the ordinal outcome distribution for a different originating class. If we treat each row X_i of our data matrix X (data on transitions from a single originating class i) as arising from observations of independent and identically distributed ordinal random variables, then these observations share a common value of μ_k in (1) denoted μ_i to acknowledge its dependence on the originating class i. Similarly, these ordinal outcomes share common threshold parameters, $\alpha_{i1}, \ldots, \alpha_{iD-1}$. A general model is therefore

$$q_{ij} \equiv \sum_{k=1}^{j} p_{ik} = g(\alpha_{ij} - \mu_i) \tag{2}$$

so q_{ij} are the cumulative transition probabilities for row i, and we assume the same form of underlying latent distribution (link function g) for each row. In fact, this is not a restriction, as any transition matrix can be fitted exactly by (2) whatever the specification for g. However, the fidelity of a more parsimonious specification for $\alpha_{ij} - \mu_i$ in (2) will depend on the form of g. In this paper, we consider two simplifications of (2), the standard cumulative link specification, introduced in (1), which for a transition matrix can be written as

$$q_{ij} = g(\alpha_j - \mu_i) \tag{3}$$

and the scale-varying cumulative link model

$$q_{ij} = g\left(\frac{\alpha_j - \mu_i}{\sigma_i}\right). \tag{4}$$

The standard cumulative link model (3) assumes a set of common thresholds, α, with the difference between rows of the transition matrix being represented by a shift (mean-change) in the distribution of the underlying latent variable. The scale-varying model allows for a ‘shift and stretch’ with the latent distribution differing between rows in both location (μ) and dispersion (σ). Note