Modelling the profitability of credit cards by Markov decision process

Mee Chi SO [a,*,1]   Lyn C THOMAS [a,2]

[a]*School of Management, University of Southampton*

Highfield, Southampton, United Kingdom, SO17 1BJ

Abstract

This paper derives a model for the profitability of credit cards, which allow lenders to find the optimal dynamic credit limit policy. The model is a Markov decision process, where the states of the system are based on the borrower's behavioural score and the decisions are what credit limit to give the borrower each period. In determining the Markov chain which best describes the borrower's performance second order as well as first order Markov chains are considered and estimation procedures that deal with the low default levels that may exist in the data are considered. A case study is used to show how the optimal credit limit can be derived.

Key words: OR in banking, Markov decision process, Credit card, Profitability, Probability of Default

_____

* Corresponding author

[1] Tel:+44 (0)23 8059 3557/Fax:+44 (0)23 8059 3844/e-mail:mekoso@soton.ac.uk

[2] Tel:+44 (0)23 8059 7718/Fax:+44 (0)23 8059 3844/e-mail:L.Thomas@soton.ac.uk

1 Introduction

Since the advents of credit cards in the 1960s, lenders have used credit scoring, both application and behavioural scoring to monitor and control default risk. However in the last decade their objective has changed from minimising default rates to maximising profit. Lenders have recognized that operating decisions are crucial in determining how much profit is achieved from a card. This paper focuses on the most important decision in an operating policy: the management of credit limit. Soman and Cheema (2002) conducted a study on the use of credit limit policies in encouraging spending and found that the availability of additional credit does promote card usage in some consumers. Consumers assumed lenders have some sophisticated models, which was used to determine appropriate credit limits, but that is not the case in reality.

So how do lenders currently decide on what credit limit to offer a credit card customer? Most use subjective policies based on a risk/return matrix, i.e. they agree credit limits for each combination of risk band and average balance, which is considered a surrogate for the return to the lender from that customer. This approach is static in that it does not consider whether or how the customers default risk and profitability to the lender will change over time. Nor is there any model to guide what are the optimal credit limits to choose.

We therefore propose using Markov Decision Processes (MDP) to improve the credit limit decision. A MDP model provides a way of making sequential decision by considering the evolution of a customer's behaviour over time. It also allows one to calculate the profitability of a credit card customer under the optimal dynamic credit limit policy. Lenders keep a wealth of historical credit card data, in particular the monthly values of a customer's behavioural score which is their way of assessing the default risk of the customer in the next year. Building the Markov decision process model on behavioural scores has the advantage that most lenders have been keeping this data on customers

for a number of years. With the advent of the Basel Accord in 2008, lenders are required to keep such data for five years and are encouraged to keep it through a whole economic cycle.

MDPs have been used in a number of different contexts (Heyman and Sobel, 1982; Ross, 1983; White, 1985, 1988, 1993; Kijma, 1997). The first application of MDPs in consumer credit was by Bierman and Hausman (1970) who looked at the repayment of a loan where no further borrowing was allowed. The model assumed the repayment of the customer followed a prior probability distribution. Using a Bayesian approach, the model revised the probability of repayment in the light of the collection history. Modifications of the basic model were made both in the accounting rules (Dirickx and Wakeman, 1976) and in the form of the Markov chain (Frydman et al., 1985) followed. The use of MDP models to manage the characteristics of consumer lifetime value can be found in Trench et al. (2003) where MDP models were used with the objective of adjusting a consumer's credit card limit or annual percentage rate (APR). The objectives in that paper are similar to the one in this paper. However their state space did not involve behavioural scores nor were they concerned with the problems that occur in estimating the transition probabilities if there are low default rates. Instead they used a six dimensional state space each dimension having only two or three categories describing the recency and frequency of purchases and payments. They developed mechanisms for reducing the size of the transition matrix through merging states. Ching et al. (2004) used MDPs to manage the customer lifetime value generated from telecommunication customers. The state space in that study again used marketing measures not risk measures and the decision was whether to implement promotions.

This paper is the first to use behavioural score bands as the basis for MDP models. The advantage of basing the model on behavioural score are considerable. Almost all lenders calculate such scores every month for every individual

both as a basis for their Basel Accord probability of default calculations, and as a way of segmenting the population on risk - see our previous discussion on risk/reward matrices.

When modelling real problems using Markov Decision Processes, the curse of dimensionality (Puterman, 1994) can mean the state space is very large and that one would need a large amount of data to obtain robust estimates of the transition probabilities. Using behavioural scores helps to overcomes this first difficulty because it itself is a "sufficient statistic" of the risk of the account and already contains information from a number of different characteristics. Also by aggregating states one can obtain a simple but meaningful state space. In our case we make part of each state an interval of behavioural scores and similarly combine possible credit limits into bands, to make up the other part of the state. We also take each of these credit limit bands to be one of the possible actions that can be chosen.

The difficulty with the quantity of data needed to calculate robust estimators of the transition probabilities is less severe in the consumer credit context because of the size of the data sets available to lenders. The only problem is that with some portfolios of loans, the number of movements directly into default from some states is so low (quite possibly zero) that the resultant estimates of zero transition probability of default may affect the structure of the Markov chain, making it non-robust. This problem of estimating default probabilities in low default portfolios also occurs in the Basel Accord mentioned earlier. We therefore use an approach suggested in that context by Pluto and Tasche (2006) and extended by Benjamin et al. (2006) which ensures the resulting Markov chain model is robust and conservative. The conservativeness is reasonable as one would prefer the model to underestimate rather than over estimate the profitability of a credit card account.

The main contribution of this paper is to show how one can use Markov

4

decision process models based on states consisting of behavioural score bands - scores which most lenders calculate on a monthly basis - to determine optimal credit limit policies in terms of profitability. The rest of the paper is organized as follows: Section 2 describes the MDP model formulation. Section 3 discusses the estimation of the transition probabilities including the probabilities of defaulting immediately. Section 4 presents the practical issues in applying the MDP model to the real credit card data and the results of the case study. The final section draws some conclusion on the model and the resultant case study.

2 The MDP model

Consider a discrete state, discrete time discounted Markov Decision Process with decision epochs $\mathcal{T}$ (indexed by $t = 1, 2, \ldots, T$) based on a state space S. Each state in the state space consists of two parts- which behavioural score band the borrower is in and what is the borrower's current credit limit band. The state space thus consists of the current credit limit band represent by $\mathcal{L}$ (indexed by $l = 0, 1, \ldots, L$) and the current behavioural score band $\mathcal{I}$ (indexed by $i = 0, 1, \ldots, I$). In our model the actions are limited to keeping the credit limit as is this period or raising it to a higher limit band. This policy of not decreasing credit limits is used by many lenders but the methodology we will describe will not change if this restriction is dropped. Thus with this limitation the action set is defined as $A_l = \{l' : l \leq l'\}$.

Two further elements need to be defined to complete the Markov decision process model. Let $p(i'|l, i)$ be the probability that if l is the current customer's credit limit band and the customer is in behavioural score band i, then the next period the customer will be in behavioural score band i'. Secondly let $r(l, i)$ be the profit obtained in the current period from a customer with credit limit l who is in behavioural score band i.

5

The objective is to maximise the discounted profit obtained from the customer over the next t periods where the discount factor λ describes the time value of money. This leads to the following optimality equation for $V_t(l, i)$, the maximum expected profit over the next t periods that can be obtained from an account which is currently in behaviour score band i, and with a credit limit of l:

$$V_t(l, i) = \max_{l' \in \mathcal{L}} \{ r(l, i) + \sum_{i'} p(i'|l, i) \lambda V_{t-1}(l', i')] \} \tag{1}$$

The right-hand-side of (1) corresponds to the profit over the next t periods if we change the credit limit to l' from l at the end of the current period for an account with behavioural score state i. We assume it takes one time period for the borrower to become aware of a change in the credit limit as this is usually included in the monthly balance statement sent to the customer. Removing this delay makes no difference to the methodology though of course the optimality equation will be slightly different. The profit to the lender from the credit card at the end of the current period is $r(l, i)$. The $p(i'|l, i)$ is the probability that the behavioural score changes to band i'. In that case, the profit on the remaining $t - 1$ period is $V_{t-1}(l', i')$. The discount factor λ is introduced because the subsequent profits in the remaining $t - 1$ periods actually occur one period after those used in calculating $V_{t-1}(l', i')$, since that assumes the $t - 1$ periods start now.

The optimality principle says that the optimal decision l', is the one that maximizes this sum of the future profit, where credit limits can only remain the same or be increased.

3   Estimating the probability of Default

Maximum likelihood estimators are usually used to estimate the transition probabilities of a Markov chain. In the Markov chain described in Section 2,

let $n_t(l,i)$ be the number of accounts in state (l,i) at time t and let $n_t^j(l,i)$ of them move to behavioural score state j at time $t+1$. Assuming the Markov chain is stationary means the maximum likelihood estimate $\tilde{p}(j|l,i)$ for the probability $p(j|l,i)$ is

$$\frac{\sum_{t=1}^{T-1} n_t^j(l,i)}{\sum_{t=1}^{T-1} n_t(l,i)}$$

In reality, moving directly to the default state is a rare event, particularly for high value (low risk) behavioural scores. There may be no examples in the data of transaction from certain states (l,i) to the default state D. Thus it is possible that $p(D|l,i)$ may be very small or even equal to zero. Putting such estimates into the MDP model leads to apparent "structural zeros" which change the connectedness of the dynamics in the state space. If the probability of defaulting from a given state is zero this can lead to unusual optimal policies because the system wants to move to those apparently "safe" states. For example, suppose there are only three behavioural states: Excellent, Good and Bad where Bad is the default state. If there are two credit limit states, 10000 and 50000.

[Figure 1 insert about here]

Provided the discount value λ is close to 1 and t is very large, the optimal policy will have a credit limit of $10,000$ in state 1 and 2 even though the profit per period $r(l,i)$ is much higher when the credit limit is $50,000$ than when it is $l = 10,000$. This is because there is apparently no chance of default if the credit limit is $10,000$ while there is a very small chance of default when the credit limit is $l = 50,000$. This anomaly arises because there have only been 800 cases when the credit limit is 10,000, but 8,000 cases when the credit limit is 50,000. Let the maximum likelihood estimates are $p(D|10000, L) = 0$ and $p(D|50000, i) \geq 0$.

One way to overcome this problem is to take a conservative estimate of the

default probabilities, rather than the maximum likelihood estimate. This problem has been extensively discussed in the context of the Basel Accord where again bank regulators and lenders have been considering the robustness of estimates of default probabilities in low default portfolios.

We will follow the approach introduced by Pluto and Tasche (2006) and extended by Benjamin et al. (2006). Firstly we assume the transitions to default are monotonically decreasing as the behavioural score increase and so if the score bands are labelled with I being the highest quality, so

$$p(D|l, I) \leq p(D|l, I-1) \leq \ldots \leq p(D|l, 2) \leq p(D|l, 1) \tag{2}$$

So a conservative assumption would be that

$$p(D|l, I) = p(D|l, I-1) = \ldots = p(D|l, 2) = p(D|l, 1) \tag{3}$$

where 1 is the most risky of the low default portfolios. This means that to estimate $p(D|l, I)$, we take a sample of $N(l, I) = \sum_{t=1}^{T-1} \sum_{i=1}^{I} n_t(l, i)$ cases when customers are in state I and $D(l, I) = \sum_{t=1}^{T-1} \sum_{i=1}^{I} n_t^D(l, i)$ of these customers are defaulted in the next month.

The second conservative assumption in this approach is not to use the MLE estimate of the default probability, but rather take the lower confidence limit of the default probability. So as there are $D(l, I)$ accounts defaulting in the next period from the $N(l, I)$ accounts under consideration. It is assumed this is given by a Binomial distribution $B(N(l, I), p)$. One choses p to be the highest probability of default, so that the corresponding lower α-confidence limit is exactly $D(l, i)$, i.e. getting a lower number of default this $D(l, i)$ has a no more than $1 - \alpha/2$ probability of occurring, since the mean and variance of the Binomial distribution is $N(l, I)p$ and $N(l, I)p(1-p)$. One choses p to be the value where

$$N(l, I)p - \Phi^{-1}(1 - \alpha/2)\sqrt{N(l, I)p(1-p)} = D(l, I)$$

where $\Phi^{-1}(1-\alpha/2)$ is the converse cumulative standard normal distribution, i.e. how many standard deviations below the mean on a standard normal distribution minus the chance of getting a lower value is $1-\alpha/2$.

One choses the estimate $\hat{p}^D(l, I)$ in this way for these states (l, I) where the number of actual defaults $\sum_{t=1}^{T} n_t^D(l, I)$ is at or below same agreed value, - which might be zero. One would like to use MLE to obtain the estimates of the other transition probabilities $\hat{p}^j(l, I)$ from state (l, I). However this would result in the sum of the transition probabilities being greater than 1 and so instead one defines there are $\alpha(l, I)\hat{p}^j(l, I)$ where

$$\alpha(l, I) = \frac{1 - \hat{p}^D(l, I)}{\sum_{j \neq D} \hat{p}^j(l, I)}$$

For states (l, i) where the number of defaults $\sum_{t=1}^{T} n_t^D(l, i)$ exceed the low default bound, MLE are used to estimate all transition probabilities.

4   Apply MDP model to credit card data

4.1   Sampling and data preparation

The MDP model developed above was applied to credit card data from a major Hong Kong bank. The dataset consisted of the credit card histories and characteristics of over 1,400,000 credit accounts for each of 60 months. The fields used in this study were account balance, account repayment, monthly profit on that account, credit limit, account written-off record and behavioural score.

In each monthly dataset, we extracted a random sample of 50,000 accounts. We looked at how these accounts performed in the next month including the change in behavioural score. There were 3,000,000 cases (50,000×60 months) extracted from the dataset. We deleted accounts with ambiguous, missing or

9

special data, which accounted for less than 0.2% of the sample and we used 2,994,602 transitions for the analysis.

One possible concern is whether the existing credit limit policy was affecting the data. However, looking at the sample, only 8% of the 3 million transitions had credit limit adjustment in the previous three months and most of these were to Inactive accounts. Thus we felt the credit limit change due to the existing card lender's policy would not have a major impact on the data.

4.2 Special accounts

Each month an account is given a behavioural score or put in a special state, such as closed, inactive, 3+Cycle or defaulted. A Closed account is one where the credit card service terminated with zero account balance. A credit card account which has never been activated or was newly opened in the last two months before the sample point (and so does not have enough data to merit a behavioural score) or has not been used in the last twelve months is called Inactive. A 3+ Cycle account is one in which the account has been in arrear for 3 or more months but the lender has not yet written the account off.

The most important special account states in the data are the defaulted or written off account. There are four possible reasons to write-off an account, bankruptcy, charge-off, revoked and 3+ cycle delinquent. Bankruptcy is when the borrower is declared bankruptcy by a court; charge-off is when the lender does not believe the debt can now be recovered by standard methods; and revoked is when the credit card is stopped because of illegal behavioural by the borrower, such as being over the credit limit persistently. Lenders pass a written off account to the debt collection department to follow-up. Such written-off account may repay all, part or none of the outstanding debt. Even when the account makes full repayment, the time of the collection process is

10

uncertain and could be several years. It is important to estimate the average future repayment amount of different default accounts because it changes the values of the profit function in (1). There is very little research (Matuszyk et al., 2007) on estimating the loss given default of revolving credit products. So we use a simple approach to compare the debt repayment ratios. Define $R \equiv \frac{B_{t+1} - B_{t+24}}{B_{t+1}}$ where B_{t+1} is the current balance of an account at default in month t. The repayment ratio R thus is defined as the proportion of the debt repaid to the lender after two years. For confidentiality reason, we can not show the exact repayment ratio, but the results showed one of the forms of default had a high repayment ratio which was significantly different from the others. We call this default account state, Bad2 and group the rest of the three forms of written off together into one default state, called Bad1. So in our cost function the loss generated by accounts in Bad1 is higher than those of Bad2.

4.3   Coarse-classifying

Since behavioural score typically has several hundreds of values, it is sensible to split it into a number of bands to reduce the size of the state space. We aim to find suitable splits by trying to get the Markovian assumption to hold as nearly as possible. To check whether the chain satisfies this assumption, for every state, we investigated the hypothesis that the probability of moving from $s_t = (l_t, i_t)$ to i_{t+1} is independent of the state at $t - 1$, i.e. $s_{t-1} = (l_{t-1}, i_{t-1})$. Define $n_t(l_{t-1}, i_{t-1}; l_t, i_t; i_{t+1})$ to be the number of times that a credit account was in state (l_{t-1}, i_{t-1}) at time $t - 1$ followed by moving to (l_t, i_t) at time t and i_{t+1} at time $t + 1$. Similarly define $n_t(l_t, i_t; i_{t+1})$ to be the number of times that a customer was in state (l_t, i_t) at time t then moving to behaviour score i_{t+1} at time $t + 1$. We assume the chain is stationary, thus the estimator for $p(i_{t+1} | l_{t-1}, i_{t-1}, l_t, i_t)$ is:

$$\hat{p}(i_{t+1}|l_{t-1}, i_{t-1}, l_t, i_t) = \frac{\sum\limits_{t=0}^{T-2} n_t(l_{t-1}, i_{t-1}; l_t, i_t; i_{t+1})}{\sum\limits_{t=0}^{T-2} n_t(l_{t-1}, i_{t-1}; l_t, i_t)} \tag{4}$$

The Markovity of the chain corresponds to the hypothesis that $p(i_{t+1}|1, 1, l_t, i_t)$ = $p(i_{t+1}|2, 1, l_t, i_t)$ = ... = $p(i_{t+1}|L, 1, l_t, i_t)$ = $p(i_{t+1}|1, 2, l_t, i_t)$ = ... = $p(i_{t+1}|L, I, l_t, i_t)$, for l_t, i_t, i_{t+1}. To check on the Markovity of state (l_t, i_t), we use the chi-square test (Anderson and Goodman, 1957). Let

$$\chi^2_{(l_t, i_t)} = \sum_{(l_{t-1}, i_{t-1}) \in S} \sum_{i_{t+1} \in I} \frac{n^*(l_{t-1}, i_{t-1}; l_t, i_t)[\hat{p}(i_{t+1}|l_{t-1}, i_{t-1}, l_t, i_t) - \hat{p}(i_{t+1}|l_t, i_t)]^2}{\hat{p}(i_{t+1}|l_t, i_t)} \tag{5}$$

where

$$\hat{p}(i_{t+1}|l_t, i_t) = \frac{\sum\limits_{t=1}^{T-1} n_t(l_t, i_t; i_{t+1})}{\sum\limits_{t=1}^{T-1} n_t(l_t, i_t)} \tag{6}$$

and

$$n^*(l_{t-1}, i_{t-1}; l_t, i_t) = \sum_{t=1}^{T-1} n_t(l_{t-1}, i_{t-1}; l_t, i_t) \tag{7}$$

Anderson and Goodman (1957) showed that (5) has a chi-square distribution with $(I - 1)(L - 1)^2$ degree of freedom.

A traditional approach is to start with a fine classification i.e. with more bands then one really wants and then check if one can combine adjacent bands. Alternatively, one can use the classification tree approach of finding the best split into two classes and then splitting one of these into two more until it is not worth splitting further. That is the approach we use in the following algorithm. We repeated the splitting process in both behavioural score and credit limit then came up with 15 score bands, of which 10 were behavioural score band and 5 were special states; we also ended up with 11 credit limit bands as listed in Table 1 and 2 respectively. For confidential reason, we do

not disclose the precise behavioural score bands. We use Score1 to Score10 to represent the credit score where Score1 are those with lowest behavioural score and Score10 are those with highest. Similarly, we use Limit1 to Limit10 to represent the credit limit with Limit1 as the lowest credit limit band. Tables 1 and 2 give the distribution of the borrowers into the different behavioural score bins and the various credit limit ranges in a typical month.

[Table 1 insert about here]

[Table 2 insert about here]

4.4   Choice of Order

A MDP is mth-order if these transition probabilities depend on which state the system is currently in and was in for the previous $m-1$ periods. For a finite order Markov chain the transition probability depends only on the current state where for a mth-order Markov chain the transition at time t depends on the states $(i_t, i_{t-1}, \ldots, i_{t+1-m})$ that it occupied for the last m time periods. So, the number of states increase exponentially in m as there are $|S|^m$ states in a mth-order MDP. To test whether a chain satisfies the mth-order Markovity assumption, one can use the chi-square test which is also used to check the homogeneity of a contingency table (Anderson and Goodman, 1957) (as will be shown in the next section). Test results indicated the Markov chains are not first order. In reality, almost all applications fail to satisfy the first-order Markovity assumption. This is because with so much data, one usually can improve the fit beyond what are the narrow significance limits. What is more important is whether there is a significant improvement in the fit, when one uses second or third order Markov chains. So, we tested whether the process is second-order Markov i.e. we redefined the state so that it carried the history of $t-2$ and $t-1$. Although there was an improvement on the chi-

13

square values, the hypothesis that the chain was a second-order MDP was also not justified in Anderson's Goodness-of-fit test. Using an even higher order Markov chain increases the size of the state space exponentially and so will affect the robustness of the model. So it is a trade off between improvements of fit and increase in size of model. Like many authors (White, 1993) We found the improvement when going to second order or higher order chains is not sufficient to warrant the loss in robustness and simplicity. Therefore, we chose to use first-order to simplify the state space as well as reducing the computational time.

4.5   Transition matrices and profit function

Table 3 and 4 show the transition matrices in the dataset. Each entry represents the transition percentage from the state of the account at time t to the state of the account at time $t + 1$. "-" represents there was no observation in the sampling period; "0" that the transition probability was less than 0.0005. The numbers of account in the different states at t are given in the last column of the table. As the transition matrix is enormous, we only presented those of Limit1 and Limit10 for illustration. Also, we show the results by splitting the states into two groups- ordinary states are those with Score2 or above; Closed, Bad1, Bad2, 3+ cycle, Inactive and Score1 are classified as special states.

[Table 3 insert about here]

[Table 4 insert about here]

Table 3 shows the transition of moving to ordinary states. The matrix is dominated by the diagonal entries as one expects. The volatility of score transition decreases as credit limit increase. 88.6% of accounts with Score10 and Limit10, remains in the same state after one month whereas it is only 75.2% of accounts in Score 10 but with Limit1 who stay in Score 10 band. This is repeated in

14

most high behavioural score groups where for low score group, the account movement are much more similar for different credit limits.

[Table 4 insert about here]

Table 4 shows the transitions of moving to special accounts. For Limit10, there is no example of an account with Score6 or above moving to the Bad1 state. However 52.8% of those in Score1, with Limit10, move to Bad1 status the next month, which is a much higher percentage then there with credit limits in Limit1. This may be due to the higher balance that such borrowers are carrying. On the other hand, as Table 3 shows 48.4% of these in Score1 and Limit1 improve their scores to Score2 in the next month.

4.6   Transition of LDP

[Table 5 insert about here]

[Table 6 insert about here]

To prevent introducing invalid structural zeros into the transition matrix, we adjust the transition matrix on a low default portfolio by the method presented in Section 3. It is important to determine when such an adjustment needs to be made. Benjamin et al. (2006) tested several approaches to determining when this adjustment is useful. They concluded that there is no ideal method but a simple approach of using it when band has below 20 cases of moving directly to default in the whole sample in which the model is built. We found there is a threshold between Score3 and Score4 in the dataset where, in most of the credit limit bands, the number of defaults in Score4 or higher score bands is below 20. We thus defined the accounts with Score4 to Score10 as the low default portfolios and the corresponding adjusted transition probabilities are presented in Table 5 and 6.

4.7 Profit Function

[Table 7 insert about here]

Table 7 shows the profit function used in the MDP model which was estimated by using the monthly profit field. For each state (l, i), the average monthly profit of the borrowers was taken over the cases when the borrower is in behavioural state i with credit limit l being applied. The profit is decreasing with behavioural score because a borrower with high behavioural score is more likely to make full repayments. Thus the lenders can only gain the merchant fee from them (the fee on each credit card purchase which is paid by the retailer to the credit card company). On the other hand, the low behavioural score accounts are more likely to accumulate debt in their accounts. Thus these borrowers generated both interest and interchange fees for the lenders. A second observation from Table 7 is that profit increases with credit limit. This is expected as the current purchased goes up is the credit limit goes up and hence so does the merchant fee.

4.8 Optimal Policy

We implemented the value iteration algorithm (Puterman, 1994) to obtain the optimal policies, and we used $\lambda = 0.995$ (a rough estimate of 6% yearly inflation rate) as our monthly discount value.

[Table 8 insert about here]

[Table 9 insert about here]

Table 8 and 9 show the optimal policies/values of using simple PDs and those of using adjusted PDs on the LDPs respectively. For example, in Table 8, the cell in row 4 and column 5 corresponds to accounts with Score2 and Limit4.

16

The optimal policy is to increase their credit limit to Limit5 and if so, in long run, the optimal gained from them is HK\$32,781.

For 3+ Cycle accounts, the optimal policies are for their credit limit to be unchanged. This follows because many of them move to default at $t+1$. Since the expected loss generated from these accounts increases with credit limit, in order to minimize the potential loss, the model chooses to keep their existing credit limit unchanged. The model suggests increasing the credit limit of some Inactive accounts to give encouragement to these borrowers to start using the card. Since there is no history of repayment the optimal policy does not move these inactive accounts to the highest credit limit band. The accounts in Score1 generate loss and are likely default in the next month. Therefore the best policy is to keep their credit limit unchanged. The optimal policies suggest increasing the credit limit of accounts in Score3 to Score5 to the highest credit limit band. These accounts are less risky and very profitable. In this model the policy is aggressive and encourages spending.

One surprise is that the optimal policies do not increase the credit limit of the highest behavioural scores, i.e. those with behavioural Score7 or above. Instead, the optimal policy is to keep their credit limit unchanged. This is because they are less profitable then those with lower behavioural score. Also, high credit limits have a relatively more severe impact on their chance of moving to risky or even default states, than an account with lower behavioural score.

5 Conclusion

This paper identifies the use of the MDP model to generate a dynamic credit limit policy. It explains how one can link the profitability of a borrower to their default risk by employing a risk estimate-the behavioural score-that is used by

every lender. The paper has also considered a conservative alternative to using MLE estimate for all transition probabilities which would avoid difficulties in low default portfolios. The effectiveness of using this method has been shown by a real credit card dataset.

References

Anderson, T. W., Goodman, L. A., 1957. Statistical inference about markov chains. Ann. Math Statist. 28, 89–109.

Benjamin, N., Cathcart, A., Ryan, K., 2006. Low default portfolios: a proposal for conservative estimation of default probabilities. Working paper, Financial Services Authority.

Bierman, H., Hausman, W. H., 1970. The credit granting decision. Management Science 16, 519–532.

Ching, W. K., Ng, M. K., Wong, K. K., Altman, E., 2004. Customer lifetime value: stochastic optimization approach. Journal of Operational Research Society 55, 860–868.

Dirickx, Y. M. I., Wakeman, L., 1976. An extension of the bierman-hausman model for credit granting. Management Science 22, 1229–1237.

Frydman, H., Kallberg, J. G., Kao, D. L., 1985. Testing the adequacy of markov chains and mover-stayer models as representations of credit behaviour. Operations Research 33, 1203–1214.

Heyman, D. P., Sobel, M. I., 1982. Stochastic Models in Operations Research Vol. 1: Stochastic Processes and Operating Characteristics. McGraw-Hill, New York.

Kijma, M., 1997. Markov Processes for Stochastic Modelling. Chapman & Hall, London.

Matuszyk, A., Mues, C., Thomas, L. C., 2007. Modelling lgd for unsecured personal loans: Decision tree approach. Discussion papers in Centre for Operational Research, Management Science and information Systems,

CORMISIS-0707.

Pluto, K., Tasche, D., 2006. The Basel II risk parameters. Springer, Ch. Estimating probabilities of default for low default portfolios, pp. 79–103.

Puterman, M., 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley, New York.

Ross, S. M., 1983. Introduction to Stochastic Dynamic Programming. Academic Press, New York.

Soman, D., Cheema, A., 2002. The effect of credit on spending decisions: the role of the credit limit and credibility. Marketing Science 21, 32–53.

Trench, M. S., Pederson, S. P., Lau, E. T., Ma, L., Wang, H., Nair, S. K., 2003. Managing credit lines and prices for bank one credit cards. Interfaces 33, 4–21.

White, D. J., 1985. Real applications of markov decision processes. Interfaces 15, 73–78.

White, D. J., 1988. Further real applications of markov decision processes. Interfaces 18, 55–61.

White, D. J., 1993. A survey of applications of markov decision process. Journal of Operational Research Society 44, 1073–1096.
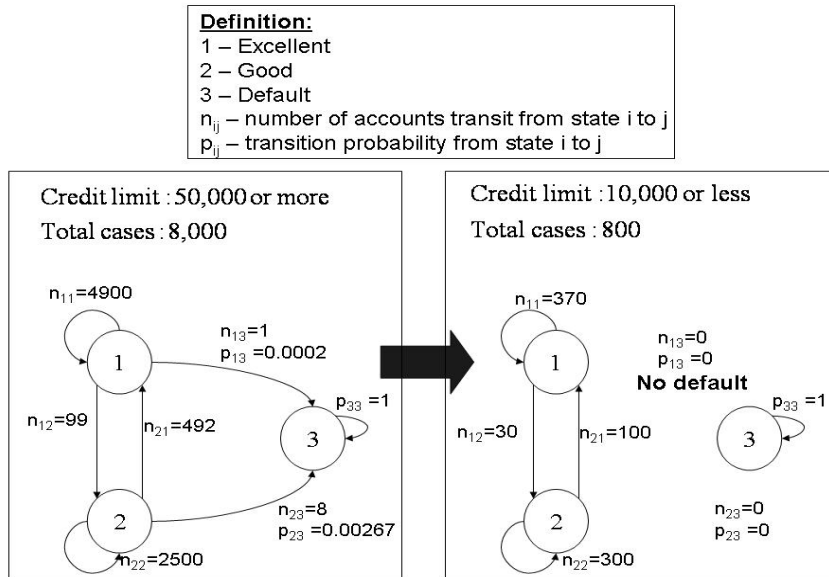
Fig. 1. An unexpected optimal policy

Index	Definition	Percentage
0	Closed	0.54%
1	Bad1	0.09%)
2	Bad2	0.05%
3	3+ Cycle Delinquent	0.02%
4	Inactive	6.16%
5	Score1	0.11%
6	Score2	1.68%
7	Score3	11.29%
8	Score4	19.30%
9	Score5	10.65%
10	Score6	5.32%
11	Score7	6.07%
12	Score8	11.20%
13	Score9	10.14%
14	Score10	17.39%

Table 1

Distribution of accounts's current score states

Index	Definition	Percentage at t
0	Closed	-
1	Limit1	7.21%
2	Limit2	10.58%
3	Limit3	11.67%
4	Limit4	10.55%
5	Limit5	9.35%
6	Limit6	10.46%
7	Limit7	10.02%
8	Limit8	10.21%
9	Limit9	9.36%
10	Limit10	10.60%

Table 2

Definition accounts' current credit limit states

		Score2	Score3	Score4	Score5	Score6	Score7	Score8	Score9	Score10	Total
	3+ Cycle	18.3	15.4	-	-	-	-	-	-	-	104
	Inactive	0.1	0.5	3.5	2.4	1	0.6	0.9	0.3	3	44616
	Score1	48.4	3.1	-	-	-	-	-	-	-	577
	Score2	66.1	25.3	0.3	-	-	-	-	-	-	8301
	Score3	7.6	76.8	13.2	0.4	0.2	0	0	0	-	27953
	Score4	0.3	7.7	77.1	8.9	1.4	1.5	1.3	0.1	0.1	42024
Limit1	Score5	0.1	2.4	13.4	60.2	6.5	5.7	5.7	2.3	2.2	19650
	Score6	0.1	1.8	7.3	15.5	48.7	6.8	12.7	2.2	3.3	9159
	Score7	0	1.3	4.7	8.3	6	57	10	6.9	3.3	11846
	Score8	0	0.6	3.1	2.9	6.9	9.4	57.3	9.3	9	18163
	Score9	-	0.4	3.9	1.8	1.3	6.2	14.4	58	13	12255
	Score10	0	0.3	1.6	0.7	0.9	1.5	9.6	9.9	75.2	21328
	3+ Cycle	13.3	13.3	-	-	-	-	-	-	-	15
	Inactive	-	0.2	2.1	2.5	1.1	0.7	1.5	0.8	2.2	25885
	Score1	24.2	3.1	1.2	-	-	-	-	-	-	161
	Score2	58.8	27.7	1.6	0.2	-	-	-	-	-	1799
	Score3	2.4	77.4	18.2	0.5	0.1	0	0	0	0	28656
	Score4	0.1	8.1	78.5	8.8	1.4	0.9	1.3	0.2	0.1	59188
Limit10	Score5	0	1.5	11.7	59.5	7.7	5.6	7.4	4.4	1.6	26563
	Score6	0	0.7	4.7	10.4	58.9	8.2	9.1	3.3	3.3	14879
	Score7	0	0.7	3.4	7	6.9	55.2	10.4	10.2	4.9	17050
	Score8	0	0.5	2.4	3.7	4.4	6.3	62.2	9.7	10.1	32561
	Score9	0	0.3	1.1	1.5	0.9	4.5	7.3	73	10.9	37615
	Score10	-	0.1	0.9	0.3	0.4	0.5	4.4	4.7	88.6	72939

Table 3

Transition matrix - moving to ordinal accounts

		Closed	Bad1	Bad2	3+ Cycle	Inactive	Score1	Total
	3+ Cycle	17.31	33.65	11.54	2.88	-	0.96	104
	Inactive	0.79	0	0	-	86.96	-	44616
	Score1	1.73	22.53	7.11	5.72	-	11.44	577
	Score2	1.07	1.31	1.06	0.48	-	4.36	8301
	Score3	0.97	0.16	0.18	0.03	0.12	0.34	27953
	Score4	1.08	0.02	0.04	-	0.39	-	42024
Limit1	Score5	1.1	-	0.04	-	0.34	-	19650
	Score6	1.19	0.01	0.05	-	0.45	-	9159
	Score7	1.18	0.01	0.03	-	1.17	-	11846
	Score8	0.91	-	0.01	-	0.5	-	18163
	Score9	0.72	-	0.01	-	0.23	-	12255
	Score10	0.33	0	-	-	0.06	-	21328
	3+ Cycle	46.67	20	6.67	-	-	-	15
	Inactive	0.89	-	-	-	88.02	-	25885
	Score1	3.73	52.8	4.35	0.62	-	9.94	161
	Score2	1.56	2.56	1.72	1.22	-	4.72	1799
	Score3	0.42	0.19	0.39	0	0.1	0.23	28656
	Score4	0.36	0.03	0.11	-	0.14	0	59188
Limit10	Score5	0.38	0	0.03	-	0.28	-	26563
	Score6	0.58	-	0.01	-	0.74	-	14879
	Score7	0.57	-	0.01	-	0.66	-	17050
	Score8	0.29	-	0.01	-	0.34	0	32561
	Score9	0.3	-	-	-	0.22	-	37615
	Score10	0.15	-	0	-	0.03	-	72939

Table 4

Transition matrix - moving to special accounts

		Score2	Score3	Score4	Score5	Score6	Score7	Score8	Score9	Score10	Total
	Score4	0.3	7.7	77.1	8.9	1.4	1.5	1.3	0.1	0.1	42024
	Score5	0.1	2.4	13.4	60.2	6.5	5.7	5.7	2.3	2.2	19650
Limit1	Score6	0.1	1.8	7.3	15.5	48.7	6.8	12.7	2.2	3.3	9159
	Score7	0	1.3	4.7	8.3	6	57	10	6.9	3.3	11846
	Score8	0	0.6	3.1	2.9	6.9	9.4	57.2	9.3	9	18163
	Score9	-	0.4	3.9	1.8	1.3	6.2	14.4	58	13	12255
	Score10	0	0.3	1.6	0.7	0.9	1.5	9.6	9.9	75.1	21328
	Score4	0.1	8.1	78.6	8.8	1.4	0.9	1.3	0.2	0.1	59188
	Score5	0	1.5	11.7	59.5	7.7	5.6	7.4	4.4	1.6	26563
Limit10	Score6	0	0.7	4.7	10.4	58.9	8.2	9.1	3.3	3.3	14879
	Score7	0	0.7	3.4	7	6.9	55.2	10.4	10.2	4.9	17050
	Score8	0	0.5	2.4	3.7	4.4	6.3	62.2	9.7	10.1	32561
	Score9	0	0.3	1.1	1.5	0.9	4.5	7.3	73	10.9	37615
	Score10	-	0.1	0.9	0.3	0.4	0.5	4.4	4.7	88.6	72939

Table 5

Transition matrix of the adjusted PDs on LDPs - moving to ordinal accounts

		Closed	Bad1	Bad2	3+ Cycle	Inactive	Score1	Total
	Score4	1.08	0.03	0.06	-	0.39	-	42024
	Score5	1.1	0.02	0.06	-	0.34	-	19650
Limit1	Score6	1.19	0.02	0.06	-	0.45	-	9159
	Score7	1.18	0.02	0.05	-	1.17	-	11846
	Score8	0.91	0.02	0.05	-	0.5	-	18163
	Score9	0.72	0.02	0.04	-	0.23	-	12255
	Score10	0.33	0.01	0.04	-	0.06	-	21328
	Score4	0.36	0.02	0.04	-	0.14	0	59188
	Score5	0.38	0.01	0.03	-	0.28	-	26563
Limit10	Score6	0.58	0.01	0.03	-	0.74	-	14879
	Score7	0.57	0.01	0.02	-	0.66	-	17050
	Score8	0.29	0.01	0.02	-	0.34	0	32561
	Score9	0.3	0.01	0.02	-	0.22	-	37615
	Score10	0.15	0.01	0.01	-	0.03	-	72939

Table 6

Transition matrix of the adjusted PDs on LDPs - moving to special accounts

Value(Policy)	Limit1	Limit2	Limit3	Limit4	Limit5	Limit6	Limit7	Limit8	Limit9	Limit10
Closed	8	25	49	95	12	41	108	33	69	302
Bad1	-4574	-8022	-11085	-16358	-17845	-25578	-34028	-41246	-47751	-88029
Bad2	193	62	286	-320	164	-355	-924	-440	-1070	-1387
3+ Cycle	-621	-647	-1068	-1040	-891	-825	-1071	-1890	-642	-248
Inactive	-8	-6	-6	-6	-2	-1	-1	-4	-6	-6
Score1	-706	-1131	-1581	-2067	-2047	-2576	-3316	-3845	-5461	-10209
Score2	202	255	367	489	555	697	894	1210	1351	2052
Score3	151	189	282	395	220	462	599	704	915	1674
Score4	33	42	80	132	77	170	235	296	429	1126
Score5	8	6	18	29	27	42	60	78	123	311
Score6	8	2	8	14	16	23	32	39	66	194
Score7	0	-2	3	5	14	16	24	36	61	142
Score8	-6	-6	-3	0	9	10	17	24	44	102
Score9	-7	-8	-5	-3	4	6	11	21	35	85
Score10	-9	-7	-6	-4	4	5	8	17	26	70

Table 7

Profit Function (in HK$)

Value(Policy)	Limit1	Limit2	Limit3	Limit4	Limit5	Limit6	Limit7	Limit8	Limit9	Limit10
3+ Cycle	10874(1)	9522(2)	5745(3)	8932(4)	5497(5)	10924(6)	12820(7)	14385(8)	3856(9)	-8693(10)
Inactive	35889(2)	36042(2)	35810(3)	35723(4)	35280(6)	35477(6)	35015(7)	34887(8)	34000(10)	33977(10)
Score1	18726(1)	17209(2)	13135(3)	6775(4)	11809(5)	2383(6)	-578(7)	-8533(8)	-16155(9)	-54220(10)
Score2	34117(2)	34303(2)	34296(3)	32781(5)	32606(5)	32117(7)	32606(7)	32742(8)	30080(9)	24065(10)
Score3	41219(10)	41834(10)	42101(10)	42348(10)	42334(10)	42469(10)	42651(10)	42889(10)	43040(10)	43956(10)
Score4	41999(10)	42216(10)	42314(10)	42426(10)	41914(10)	42275(10)	42369(10)	42309(10)	42578(10)	43461(10)
Score5	38371(10)	38540(10)	38626(10)	38614(10)	38355(10)	38474(10)	38546(10)	38482(10)	38530(10)	38744(10)
Score6	37191(10)	37386(10)	37340(10)	37340(10)	37232(10)	37219(10)	37273(10)	37169(10)	37225(10)	37343(10)
Score7	36618(2)	36925(2)	36878(3)	36810(10)	36738(10)	36773(10)	36732(10)	36668(10)	36731(10)	36744(10)
Score8	36242(2)	36493(2)	36353(10)	36363(10)	36403(10)	36374(10)	36370(10)	36381(10)	36373(10)	36418(10)
Score9	36057(2)	36131(2)	36027(3)	35966(4)	35804(10)	35858(10)	35831(10)	35754(10)	35799(10)	35827(10)
Score10	35751(2)	35789(2)	35614(3)	35498(4)	35447(10)	35452(10)	35428(10)	35410(10)	35403(10)	35449(10)

Table 8

Optimal Policy (Using the transition matrix with no adjustment on PDs)

Value(Policy)	Limit1	Limit2	Limit3	Limit4	Limit5	Limit6	Limit7	Limit8	Limit9	Limit10
3+ Cycle	10047(1)	8731(2)	5062(3)	8138(4)	4698(5)	9976(6)	11497(7)	12893(8)	2937(9)	-9310(10)
Inactive	33154(2)	33296(2)	33017(3)	32896(4)	32379(6)	32567(6)	32110(7)	32008(8)	31192(10)	31170(10)
Score1	17411(1)	15891(2)	11957(3)	5753(4)	10550(5)	1325(6)	-1615(7)	-9507(8)	-17071(9)	-54888(10)
Score2	31961(2)	32139(2)	32134(3)	30681(5)	30519(5)	30032(7)	30512(7)	30607(8)	28026(9)	22077(10)
Score3	38614(10)	39202(10)	39462(10)	39704(10)	39666(10)	39820(10)	40000(10)	40230(10)	40384(10)	41290(10)
Score4	39038(10)	39249(10)	39347(10)	39459(10)	38932(10)	39304(10)	39396(10)	39326(10)	39595(10)	40480(10)
Score5	35227(10)	35394(10)	35473(10)	35457(10)	35201(10)	35316(10)	35373(10)	35301(10)	35354(10)	35543(10)
Score6	34151(2)	34336(2)	34159(10)	34149(10)	34045(10)	34047(10)	34079(10)	33974(10)	34022(10)	34116(10)
Score7	33719(2)	34004(2)	33877(3)	33728(4)	33562(10)	33587(10)	33543(10)	33470(10)	33538(10)	33529(10)
Score8	33357(2)	33598(2)	33292(3)	33205(4)	33207(10)	33175(10)	33166(10)	33178(10)	33175(10)	33195(10)
Score9	33197(2)	33270(2)	33059(3)	32928(4)	32671(6)	32724(6)	32623(10)	32551(10)	32587(10)	32604(10)
Score10	32909(2)	32954(2)	32666(3)	32488(4)	32276(6)	32276(6)	32229(10)	32209(10)	32202(10)	32239(10)

Table 9

Optimal Policy (Using the transition matrix with adjusted PDs on LDPs)