

Debtor Level Collection Operations Using Bayesian Dynamic Programming

Mee Chi SO^{*(1)}, Christophe Mues⁽¹⁾, Adiel T. de Almeida Filho⁽²⁾ and Lyn C Thomas⁽¹⁾

(1) Southampton Business School, University of Southampton, Southampton, UK, SO17 1BJ

*Corresponding author/ e-mail: M.So@soton.ac.uk / phone: +44 (0)23 8059 8964

(2) Federal University of Pernambuco, UFPE, Cx. Postal 7462, CEP 50.630-970, Recife-PE,
Brazil.

Abstract

After a borrower defaults on their repayment obligations, collectors of unsecured consumer credit debt have a number of actions they can take to secure some repayment of the debt. The operations management challenge in this setting is to decide which of these actions to take, how long to take them, and in what sequence to take them, in order to maximize the recovery rate. In this paper, we adopt a dynamic programming approach to find an optimal policy of which action to undertake in the next period, using Bayesian updating to take into account the individual debtor's repayment performance thus far. The use of the model is empirically illustrated using data provided by a European bank's in-house collections department.

Keywords: Finance; dynamic programming; Bayesian updating; stochastic processes.

1. Introduction

Collectors of unsecured consumer credit debt can select from a series of different actions to secure some repayment of defaulted debt. They can include gentle actions such as reminder phone calls or polite letters. These may be followed by tougher actions involving more

persistent letters, or repeated phone calls through which the collector seeks an agreement with the borrower on a new schedule for repayment. If these actions fail, collectors could make home visits, but more often they may seek legal proceedings through the courts and the use of court officers to recover property in recompense. The collections policy is to decide: which of these actions to take; how long to take each one; and in what sequence to take them. Where the latter is concerned, it is reasonable to assume that the actions can be ordered in terms of their perceived harshness. In this paper, we assume that the collector will not be moving from a harsher action to a milder one. For example, after legal proceedings have been threatened, there seems little point in sending a gentle reminder letter. However, if this assumption were to be dropped, the proposed model could be easily modified to accommodate this.

When managing the consumer credit repayment process, lenders frequently use behavioural scoring to predict the risk of default for an individual borrower. Since input data on recent or current arrears typically feature in those models, many of the lender's decisions (e.g. whether to grant a credit limit increase) depend on the individual consumer's performance prior to default. It therefore seems natural that a lender's collection policy too should be based on the individual consumer's repayment performance while in default. For example, it could depend on how many payments the defaulter has made under the current action in the collection process, relative to the length of time the action has been undertaken. This can be more effective than simply applying the same fixed policy to everyone and may partially answer the criticism that consumer credit decisions are often too impersonal (Johnson, 2004). Also, it goes some way towards addressing regulatory concerns about responsible lending implying that collectors should be "fair in their treatment of debtors" and "establish proportional policies" (Office of Fair Trading, 2011). To achieve this, the proposed model will use a Bayesian updating approach; thus, a decision shall be made to either continue with the same action for another period, or proceed to the next action.

The success of the collection process is measured by the recovery rate, i.e., the percentage of the defaulted amount that is recovered during that process. Some institutions and regulators do not allow the collection costs to be added to the original default amount that is subject to collection, while others do. This paper assumes that one cannot collect more than the original default amount; hence, a collector seeking to maximize their profit would wish to maximize the collected amount less the collection costs, i.e.:

$$\text{Profit} = \text{Recovery Rate} \times \text{Default Amount} - \text{Collection Costs.}$$

Optimizing this for an individual defaulter is equivalent to optimizing:

$$\text{True Recovery Rate} = \text{Recovery Rate} - (\text{Collection Costs} / \text{Default Amount}).$$

We call this expression the ‘true’ recovery rate whilst referring to the first term in its right-hand side as the ‘formal’ recovery rate. The collector’s objective is to maximize the discounted true recovery rate, subject to the formal recovery rate having to remain less than or equal to 100%. A slight modification of the suggested approach would be to model the case where collection costs could instead be added to the original default amount.

The remainder of the paper is organized as follows. The following section briefly reviews the literature on recovery rates for unsecured consumer loans and the collection process involved. In section 3, we introduce the Bayesian dynamic programming model of the collection process, while section 4 describes the form of the optimal collection policy, with details of the proofs given in Appendix. Section 5 applies the model to data provided by a European bank on their in-house collection process. Next, section 6 looks at two alternative policies that are suboptimal but easier to calculate; these will give an indication of the potential gains of using the optimal Bayesian policy. Section 7 draws conclusions about the problem and the model used to solve it.

2. Literature Review

The collection process for defaulted consumer credit is a multi-billion dollar operation, yet it is little researched compared with other operations management activities. Similarly, until more recently, statistical models for estimating the recovery rate or its complement, loss given default (LGD), gained far less attention than those for the probability of default (PD) of loans. Yet, for each segment of a bank's loan portfolio, estimating LGD is a requirement for the (Advanced) Internal Ratings Based (IRB) approach under the Basel II and III regulatory frameworks (BCBS, 2004; BCBS, 2011).

Although it is common practice to use regression to model LGD, the accuracy of the loan-level estimates produced tends to remain poor, particularly so for unsecured consumer credit (Loterman et al., 2011; Zheng and Thomas, 2012), unless the strategic and operational decisions that affect recoveries could be modelled as well. Whereas for corporate credit, this process is fairly specific to the assets of the defaulted firm, the administrator's remit and whether there are potential purchasers, for consumer credit, the homogeneity of the pool of defaulted loans and the limited set of collector's actions (Coleman, 2004) allow generic models of the collection process based on the data available.

In secured consumer credit, Lucas (2006) modelled collections as a two-stage process that may involve repossession (i.e. foreclosure), followed by the sale of the security. The probability of having to repossess is first estimated using logistic regression; in the second stage, a repossession "haircut" is estimated – the percentage loss when comparing the actual sale value of the security with the expected market value. Leow et al. (2014) showed how the economy affected these estimates.

One of the first approaches to model the collection process for unsecured consumer default was by Mehta (1972) who developed a Markov process model of the cash flow behaviour under

different collection actions. The paper by Murgia and Sbrilli (2012) used neural nets, integer programming and hidden Markov chain models to score the distressed debt. It then employed a decision support system to suggest which action to be taken. However, the latter is more akin to a simple collection score approach, which does not guarantee any optimality. Matuszyk et al. (2010) suggested that decision trees are a useful way of determining the strategic decisions of when and whether to recover “in house”, use an agent or sell off the debt.

The model that is closest to this paper’s is that outlined in De Almeida Filho et al. (2010). This previous paper used a dynamic programming approach to find the optimal sequence of collection actions and the optimal duration of each action. However, an important limitation of this earlier model is that it only dealt with the problem at the portfolio level. Although this is useful for determining the staffing resources and specific expertise required to undertake the optimal strategy, it does not provide any insights as to how such strategy might vary with the individual borrower as they partially repay or fail to repay the debt. In contrast, the new model we propose in this paper works at the individual defaulter level and finds, for each of them, the collection action to be applied in the light of their repayment performance thus far.

The collection optimization method proposed in this paper adopts a Bayesian dynamic programming approach in which uncertainty about a parameter value is described by a Beta distribution. This idea was first introduced to credit granting models by Bierman and Hausman (1970), who applied it to the problem setting of a consumer who may be offered a series of one-period loans, each of which is due for repayment before the decision on the next loan is made. To develop a model that helps the lender decide whether to grant additional credit, Bierman and Hausman (1970) assumed the probability of the next loan being paid has a Beta distribution whose parameters are updated in light of the consumer’s past payment history. Waldmann (1998) further extended this model by including economic factors which may affect

the transition probability between the belief in one period and the next. In both papers, a loan is either repaid in full or not at all, and the only options the lender has are to either refuse or accept a subsequent loan. Instead, we address the problem setting of debt collections, where the objective is to maximize the recovery rate achieved. In each period, the defaulter could make a partial payment or no payment towards the default balance. Moreover, rather than having to make a simple binary decision, the collector must choose between a number of different collection actions.

The basics of the Bayesian dynamic programming approach which solves these models are outlined in Puterman (2005) and Bauerle and Rieder (2011). The latter showed how the models can be applied in finance. An alternative approach to solving partially observable dynamic programming problems uses the piecewise linearity of the value function in terms of the unknown probability. This approach was introduced by Smallwood and Sondik (1973) and a review of the different solution approaches is found in Lovejoy (1991).

3. Collections Process Model

In our proposed (infinite-horizon) dynamic programming model to determine the optimal actions for the collection process, we denote the states by (r, s, m, i) . Here, i is the current collection action; r is the ‘formal’ recovery rate that had been realized during the actions that preceded action i (as a percentage of the defaulted loan amount, D , and excluding the costs of recovery and any discounting of the recovery cash flows); s is the number of periods that action i has been undertaken; and m is the number of repayments that have occurred during those s periods under action i ($m \leq s$). Hence, as the collector starts with action i , the debt owed is still $(1 - r)D$. Now, let $F_i(m)$ be the cumulative fraction of this outstanding formal debt that is subsequently recovered through the first m payments under that action, whereas $f_i(m) =$

$F_i(m) - F_i(m - 1)$ will be used to denote the fraction received back in the m -th payment of action i .

Whether the borrower repays in any given period under action i is assumed to be a Bernoulli random variable whose probability of occurring is p_i . The collector's belief about p_i is given by a Beta distribution with parameters $B(\alpha, \beta)$, the prior conjugate distribution to a Bernoulli random variable. Under our notation, this Beta distribution has a density function $g(p)$, where $g(p) \propto p^{\alpha-1}(1-p)^{\beta-1}$; hence, its expected value is $\frac{\alpha}{\alpha + \beta}$. At the start of the collection process, the initial belief about p_i is given by the Beta distribution $B(m_0^i, s_0^i - m_0^i)$. After every successful payment period, the belief about the repayment probability changes from $B(m, s - m)$ to $B(m + 1, s - m)$, while after every period with no repayment, it changes from $B(m, s - m)$ to $B(m, s + 1 - m)$. This is a simple application of Bayes theorem. If the prior distribution is $B(1, 1)$, this corresponds to a uniform prior where the density function of p is 1 over the whole region $[0, 1]$. However, one can introduce more informative priors where $(m_0^i, s_0^i - m_0^i)$ could depend on previously collected data concerning the defaulter. This is akin to a "collection score".

As argued previously, we shall assume the collector's actions, $i = 1, \dots, I$, are ordered in increasing level of harshness. In each state (r, s, m, i) , there are thus two options available to the collector: continue with the current action i , or move to action $i+1$. This corresponds to the normal situation in collections where the position of the debtor may be reviewed every month. Since moving to a different action is immediate, note that the collector in effect also has the option of moving directly to action $i+2$ (or higher) in that same period, by undertaking action $i+1$ for 0 periods. Hence, the available actions in this state are action i and all higher actions. Now, let $V(r, s, m, i)$ be the expected 'true' future discounted recovery rate starting from that state (i.e. excluding what went on before, but allowing for the costs, c_i , incurred in each period

and discounting both the amounts recovered and the costs; each time period we discount them by a factor β). Thus, $V(r, s, m, i)$ satisfies the following dynamic programming optimality equation (Puterman, 2005):

$$V(r, s, m, i) = \max \begin{cases} \frac{(m + m_0^i)}{(s + s_0^i)} [(1 - r)f_i(m + 1) + \beta V(r, s + 1, m + 1, i)] + \left(1 - \frac{(m + m_0^i)}{(s + s_0^i)}\right) [0 + \beta V(r, s + 1, m, i)] - c_i & (1) \\ V(r + (1 - r)F_i(m), s_0^{i+1}, m_0^{i+1}, i + 1) \end{cases}$$

The first expression on the right-hand side (RHS) of the optimality equation corresponds to the decision to persist with action i for another period. Given the belief about the probability of repayment, one expects the chance of repayment is $\frac{(m + m_0^i)}{(s + s_0^i)}$. A payment is equivalent to a fraction $(1 - r)f_i(m + 1)$ of the original debt. In the event of a payment being made, the debtor will be in state $(r, s + 1, m + 1, i)$ at the start of the next period, as the posterior belief about the repayment probability then becomes $B(s_0^i + s + 1, m_0^i + m + 1)$. The chance of there being no repayment in the current period is $\left(1 - \frac{m + m_0^i}{s + s_0^i}\right)$. In that case, at the start of the next period, the borrower will be in state $(r, s + 1, m, i)$, and the belief about the repayment will change to $B(s_0^i + s + 1, m_0^i + m)$.

The second expression on the RHS of the optimality equation corresponds to the alternative decision to start action $i + 1$. This change of action is immediate (hence no β in the expression) and costless (hence no c_i in the expression). It implies that the lender starts that action with a formal recovery rate of $r + (1 - r)F_i(m)$. The probability of repayment in the next period then has a Beta distribution with parameters $(m_0^{i+1}, s_0^{i+1} - m_0^{i+1})$.

Note that, if the collection costs were to be included in the debt total that can be recovered, it would be sufficient to change the second expression on the RHS of (1) to:

$$V(r + (1 - r)F_i(m) - sc_i, s_0^{i+1}, m_0^{i+1}, i + 1).$$

Here, the first component could turn negative if the collection costs added to the debt are greater than the recoveries so far.

The standard way of solving stochastic dynamic programming problems and finding the form of the optimal policy is value iteration, in which one solves the optimality equation using the following iteration scheme:

$$V_0(r, s, m, i) \equiv 0, \forall r, s, m, i$$

$$V_{n+1}(r, s, m, i) = \max \begin{cases} \frac{(m + m_0^i)}{(s + s_0^i)} [(1 - r)f_i(m + 1) + \beta V_n(r, s + 1, m + 1, i)] + \left(1 - \frac{(m + m_0^i)}{(s + s_0^i)}\right) [0 + \beta V_n(r, s + 1, m, i)] - c_i \\ V_n(r + (1 - r)F_i(m), s_0^{i+1}, m_0^{i+1}, i + 1) \end{cases} \quad (2)$$

It is well known (Puterman, 2005) that, for compact state space problems such as (1), value iteration (2) converges to the solution of the optimality equation (1) and the optimizing actions of the RHS of (2) converge to the optimal policy. Although r is essentially a continuous variable with values between 0 and 1, we will need to discretize it in order to solve the problem numerically. In fact, this does not pose any real problem because the first action should always start with r being 0; the second action must then start with r taking a value of the form $1 - F_1(m)$, for some integer m ; similarly, the r values when the third and subsequent actions start can only take one of a finite number of values corresponding to how many repayments were made under the previous actions.

4. Form of the Optimal Policy

The form of the optimal policy is described by the two theorems of this section. Their proofs and those of their supporting lemmas can be found in appendix. The optimal policy turns out to be a control limit where, for a given r , m and i , one moves to the next action after a set number of periods, s^* . Formally:

Result 1: The optimal policy in state (r, s, m, i) is given by a set of functions, $s^*(r, m, i)$, such that:

- i) One continues with action i if $s < s^*(r, m, i)$;
- ii) One moves to action $i+1$ if $s \geq s^*(r, m, i)$.

One might expect to find that $s^*(r, m, i) \leq s^*(r, m+1, i)$; i.e., the more repayments under action i , the longer one would stick with it before swapping to another action. At first, this appears intuitive since seeing more payments under action i increases the belief that a repayment will occur in the next period. However, more payments under action i also imply a lower reward from future payments under that same action, if $F_i(\cdot)$ is concave. Hence, the two effects can cancel each other out, as a result of which the optimal action may change. If, however, the value of each repayment under action i is assumed constant, the second effect does not occur. In that case, the more payments made under action i , the longer one stays with action i , as the following result shows.

Result 2: If the recovery rate for each payment under action i is constant, i.e. $f_i(m) = f_i$, then $s^*(r, m, i) \leq s^*(r, m+1, i)$.

5. Numerical example

The data feeding into our numerical example were supplied by a European bank's in-house collections department. The dataset consisted of information on 3,084 defaulted unsecured consumer loans. For each loan, there were details of when the borrower defaulted, and how much they defaulted for. There were also three years of the subsequent collection history for each of the loans, recorded on a monthly basis. This included which collection action was applied to that debtor in each time period and how much the debtor repaid that month.

There were three actions that could be applied to a debtor. Action 1 involved communicating with the debtor to arrange and encourage an agreed repayment schedule. Action 2 involved using legal procedures to recover some of the debt. Action 3 was essentially writing off the debt in that it involved no costs but also no expectation of any further recovery of the

outstanding debt. The collection process must always begin by applying Action 1, since, in the country involved, it was a legal requirement to have discussed possible repayment schedules with a debtor before the lender can apply for legal redress. At the end of the collection process, once a debt has reached the passive write-off involved in Action 3, the lender would not return to Actions 1 or 2. The collection policy which led to this data was based on a review of all accounts undertaken once a year at the financial year end, not on the basis of any optimization strategy comparable to ours. Hence, we only use the data to define the action structure and obtain reasonable cost and recovery rate function estimates. Note that the true test of our model would be to implement the proposed policy and then observe (and if possible compare against a control sample) the resulting actual collections and costs; however, we were not in a position to do so. Instead, in order to have some indication of potential efficiency gains, later on in section 6, we will compare the expected recovery rates for the optimal policy against those obtained for two simpler reference policies.

The costs and recoveries under Action 3 are taken to be $c_3 = 0$ and $f_3(m) = 0$ for all m . For Actions 1 and 2, the monthly cost is expressed as a fraction of the debt outstanding on default. This was calculated by dividing an estimate of the total costs incurred during an action by the total number of months that that action was performed summed over all debtors. This gives the average monthly cost of that action which is then divided by the average default balance over all debtors. This leads to the values $c_1 = 0.00048656$, $c_2 = 0.00398703$.

In section 3, we proposed that the size of the debt repayments under each action i be a function of the number of previous payments under that action. For the same dataset, looking at averaged recovery data, De Almeida Filho et al. (2010) reported that an exponential function provided a good fit to the shape of these payment recovery rates. We make a similar choice in our current numerical example and define:

$$f_i(m+1) = a_i e^{-b_i m}, \text{ and hence } F_i(m) = \sum_{j=1}^{j=m} f_i(j) = \frac{a_i(1 - e^{-b_i m})}{(1 - e^{-b_i})} \quad (3)$$

Fitting to the data from Action 1 and 2 leads to: $a_1 = 0.035721$, $b_1 = 0.1470$, $a_2 = 0.0245225$, $b_2 = 0.0577$. Note that the usefulness of our model by no means depends on selecting this particular function; on different datasets, other functional forms might provide a better fit. Moreover, it is possible to segment the loan set further and fit different functions to different groups of loans, if their repayment patterns were found to vary. This would also allow introducing different cost parameters. Doing so would require solving not one but a (small) number of different model instances (which remains entirely feasible) and then applying the different resulting policies to these respective loan segments.

In our numerical example, we further assume that, at the start of each action, we have no prior knowledge concerning the probabilities of a debtor repaying. This corresponds to choosing a flat uniform prior over the interval $[0,1]$ of possible probabilities for Actions 1 and 2, before updating those probabilities based on repayment performance. In other words, we set: $m_0^1 = m_0^2 = 1$; $s_0^1 = s_0^2 = 2$ (more on alternative choices of priors in the conclusions section).

The Markov decision process (MDP), whose optimality equation is (1), can be solved using the iterative scheme shown in (2). As it stands, there are an infinite number of states in the MDP, but if we put an upper limit of S on the number of periods any specific action lasts, then, with m and s being integer values, $s \leq S$, and $m \leq s$, the number of states for Action 1 (where $r=0$ and $i=1$) becomes finite. When the decision is to proceed to Action 2, the process must move to one of the $S+1$ states $(F_1(m_1), 0, 0, 2)$, where $0 \leq m_1 \leq S$, depending on how many payments, m_1 , were received under action 1. Thus, under the second action, the number of states only multiplies with a further factor of $(S+1)$.

On solving the optimality equation, we arrive at the optimal policy, shown in Figures 1 and 2. Figure 1 shows whether the optimal decisions in the states $(0,s,m,1)$ are to stay with Action 1 or move to Action 2. The results in Figure 2a and 2b describe the decisions in states $(F_1(j),s,m,2)$ for $j=1, 10, 30$ and 50 , i.e., whether to stay with Action 2 or move to Action 3.

[Figure 1, 2a and 2b about here]

Note that keeping m fixed and increasing s (i.e. for a longer period of time and no additional payments), one moves horizontally along these graphs from left to right. In every case, we find that initially the collector should stick to the current action but then, after some s^* number of periods, move to the next action. This is the result proved in Result 1. Moving vertically up the graphs, i.e., keeping s fixed and increasing m , there are several examples of non-monotonicity, where initially the collector is told to move to the next action for low values of m (implying little chance of repayment), then stay with the existing action as m increases, and finally move again to the next action if m becomes large. The latter upper boundary on whether to change to the next action comes about because, as the number m of previous payments increases, the assumed recovery rate on the next payment, $f_i(m+1)$, eventually decreases until it becomes too low to compensate for the cost, c_i , of staying in action i . One tends to stay longer with Action 1 than with Action 2 for the simple reason that, in our numerical example, Action 1 is less costly than Action 2 to undertake.

If the debtor made more payments under Action 1, then $F_1(m)$, the cumulative recovery rate under that action, increases. Figure 2a and 2b shows that, under that scenario, there are fewer states where one opts for Action 2 before writing off the debt (cf. the shrinking green-coloured area for the states $(F_1(30),s,m,2)$ and $(F_1(50),s,m,2)$, shown on the right of Figure 2a and 2b). The reason for this is that there is less money left to recover whilst the cost of Action 2 remains fixed.

Looking at the boundary between staying with Action 1 and moving to Action 2, there are several situations where the optimal action in $(r,s,m,1)$ is Action 1, but in $(r,s+1,m+1,1)$ it is Action 2. Since a payment entails that the debtor moves from $(r,s,m,1)$ to $(r,s+1,m+1,1)$, these correspond to occasions when it is optimal to move to a new action even if the debtor has just made a payment under the current action.

[Figure 3, 4a and 4b about here]

To help us better understand the latter result, a second example considers the hypothetical case where the debtors must pay a fixed repayment amount under a given action. For example, let us use the same cost values and assume the fixed repayment rate is 0.0104 for all states. The resulting policy when in Action 1 is given by Figure 3; that in Action 2 is given in Figure 4a and 4b. Now there is only a lower boundary between both action choices as m increases. As shown by Result 2, a constant recovery rate implies that, if action i is optimal in state (r,m,s,i) , it remains optimal in $(r,m+1,s+1,i)$, as the return under action i no longer diminishes as m increases. Note that, in Figure 4a and 4b, this lower boundary between Action 2 and Action 3 moves up as more of the debt has already been collected by the first action. Therefore, the second action becomes less profitable and is used less frequently. Eventually, at $F_1(59)$, the Action 2 region disappears entirely as so much of the debt has been recovered in Action 1, that there is no point in continuing the collection procedure.

6. Bayesian myopic policy and non-Bayesian policy

The optimal Bayesian policy involves significant calculations to solve the corresponding dynamic program. Alternatively, one can obtain a near-optimal Bayesian policy using the myopic approach, which does not consider the entire time horizon but only considers one period at the time. Under this policy, one would continue to use action i in the next period, provided that it is profitable to do so in that period, i.e. if equation (4) is satisfied.

$$\left(\frac{m + m_0^i}{s + s_0^i}\right)(1 - r)f_i(m + 1) \geq c_i \quad (4)$$

As $\left(\frac{m + m_0^i}{s + s_0^i}\right)$ is the expected probability that there will be a payment in the next period, (4)

requires that the expected repayment is greater than the cost; otherwise, one moves to action $i+1$. The form of this policy (for Action 1 only), assuming the same costs and recovery functions given in section 5, is described in Figure 5. In Figure 5, there are states where the optimal policy is to move to Action 2, while the myopic policy proposes Action 1. This follows if the expected real reward under Action 1 is positive but is still less than that under Action 2. Note that we do not show the policy for Action 2 as, in this case, the result is identical to that of the optimal policy (i.e. Figure 2a and 2b) because, in this example setting, Action 3 does not incur any cost or generate any revenue.

[Figure 5 about here]

Another policy would be to ignore the Bayesian updating aspect of the policy and assume the same probability of repayment under action i , irrespective of the debtor's repayment pattern.

To ensure a like-for-like comparison, take this to be $\left(\frac{m_0^i}{s_0^i}\right)$ for action i , which amounts to 0.5

for each of the actions in the numerical example. We can derive the optimal policy in this non-Bayesian case, by solving the following dynamic programming problem:

$$V(r, m, i) = \max \begin{cases} \frac{m_0^i}{s_0^i} [(1 - r)f_i(m + 1) + \beta V(r, m + 1, i)] + \left(1 - \frac{m_0^i}{s_0^i}\right) [0 + \beta V(r, m, i)] - c_i \\ V(r + (1 - r)F_i(m), 0, i + 1) \end{cases} \quad (5)$$

Using the same costs and recovery rates, we thus obtain the optimal policy shown in Figures 6a and 6b. This policy looks very different to the Bayesian policy since the action is now independent of s . It only depends on m as the recovery rate still depends on m . For example, this policy says switch from Action 1 to Action 2 when there have been 16 payments under

Action 1 (Figure 6a). Similarly, the change from Action 2 to Action 3 depends on the number of payments made under Action 1 (Figure 6b).

[Figure 6a and 6b about here]

The value functions for the three different policies (i.e. their expected ‘true’ recovery rates) are listed in Table 1. The value reported for the model using the optimal Bayesian policy is 0.1926; i.e., it expects a recovery rate of just over 19%. The expected recovery rate for the myopic policy is 0.1894, which is a drop of just over 2%. The next line indicates that the optimal non-Bayesian policy gives an expected recovery rate of 0.1747, which is 9% lower than that of the Bayesian optimal policy. This indicates there are clear potential gains to the method proposed in sections 4 and 5 in terms of profit, under the assumptions it makes about the repayment process. Note again that these are *expected* recoveries – the *actual* recoveries if the policy were applied shall remain unknown as any variation in collection strategy would inevitably impact what returns are observed. Also, Table 1 each time reports a single number, but if a loan segmentation approach were adopted, different expected recovery rates would ensue for different loan segments.

[Table 1 about here]

7. Conclusions

This paper looked at how debt collection departments should manage their operations to maximize recoveries. It develops collection policies for defaulted consumer debt based on the performance of the individual customer. Debt collection is a multi-billion dollar industry, but also one with a high political profile in terms of legislation and media attention, much of it negative. Therefore, the collections process needs to steer a delicate course between efficiency and fairness. The actions pursued to recover a defaulted debt need to be defensible under both objectives. An optimal Bayesian collections policy would therefore seem sensible since the collector’s actions depend on the debtor’s performance.

The flexibility of the Bayesian approach is that the extra knowledge the collectors have can be incorporated in the prior distribution used. For example, the parameters m_0^i and s_0^i of the prior could be functions of the characteristics of the loan, the defaulter and the defaulter's performance before default, or even their performance under the earlier collection actions. Thus it could become an extension of the idea of a collections score. The parameters of the prior could also be made functions of the economic conditions. In that case, one could estimate the recovery rate (and thus the LGD) in downturn conditions as well as normal ones.

One can extend the model in other ways. For example, one could include a fixed cost for starting a new action by making the costs a function of s as well as i with $c_i(1) \neq c_i(s)$, $s > 1$. Moreover, one could allow for the fact that some actions have to be performed for some period, say S , until the debtor starts to repay. This would mean the recovery rate is a function of s as well as m , with $F(m, s) = 0, \forall s \leq S$. Further extensions could include having a nested action structure, in which different variations or levels of the same action can be tried; this would lead to a model formulation with one extra state variable and an additional expression in (1) for each such decision option.

Finally, in our numerical example, we found a considerable difference in the optimal policy and value between the Bayesian and non-Bayesian formulations. This suggests that there may also be significant advantages in terms of financial returns when introducing a Bayesian approach to the modelling of the collections process.

References

Basel Committee on Banking Supervision (BCBS) (2004) International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Basel Committee on Banking Supervision, Bank for International Settlements.

Basel Committee on Banking Supervision (BCBS) (2011) Basel III: a Global Regulatory Framework

- for More Resilient Banks and Banking systems. Basel Committee on Banking Supervision, Bank for International Settlements.
- Bauerie N and Rieder U (2011) *Markov Decision Processes with applications to Finance*, Springer, Heidelberg.
- Bierman H and Hausman WH (1970) The credit granting process. *Management Science* **16**: 519-532.
- Coleman AM (2004) *Collection Management Handbook*, John Wiley, Hoboken.
- De Almeida Filho AT, Mues C and Thomas LC (2010). Optimizing the Collections Process in Consumer Credit. *Production and Operations Management* **19**: 698-708.
- Johnson RW (2004). Legal, social and economic issues in implementing credit scoring in the United States. In: Thomas LC, Edelman DB, Crook JN (eds.) *Readings in Credit Scoring*, 5-15, Oxford University Press, Oxford.
- Leow M, Mues C and Thomas LC (2014) The economy and loss given default: evidence from two UK retail lending data sets, *Journal of the Operational Research Society* **65**: 363-375.
- Loterman G, Brown I, Martens D, Mues C and Baesens B (2011) Benchmarking regression algorithms for loss given default modelling. *International Journal of Forecasting* **29**: 161-170.
- Lovejoy WS (1991) A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operational Research* **28**: 47–66.
- Lucas A (2006) Basel II Problem Solving, QFRMC Workshop and Conference on Basel II & Credit Risk Modelling in Consumer Lending, University of Southampton, U.K.
- Matuszyk A, Mues C and Thomas LC (2010). Modelling LGD for unsecured personal loans: decision tree approach. *Journal of Operational Research Society* **61**:393-398.
- Mehta D (1972) Markov Process and Credit Collection Policy, *Decision Sciences* **3**: 27-43.
- Murgia G and Sbrilli S (2012) A decision support system for scoring distressed debts and planning their collection. In: *Methods for Decision making in an uncertain environment*, 69-89, Proceedings of the XVII SIGEF Conference, Reus, Spain.
- Office of Fair Trading (2011). Debt Collection, OFT guidelines for all businesses engaged in the recovery of consumer credit debt, OFT, London.
- Puterman ML (2005) *Markov Decision Processes: Discrete Stochastic Dynamic Programming (2nd ed.)*, Wiley, New Jersey.
- Smallwood RD and Sondik EJ (1973) The optimal control of partially observable Markov processes over a finite horizon. *Operations Research* **21**: 1071–1088.

Appendix

Lemma 1: $V_n(r, s, m, i)$ defined in (2) is non-decreasing in n and hence it converges to $V(r, s, m, i)$, the solution of (1).

Proof. The proof is by induction to show $V_{n+1}(r, s, m, i) - V_n(r, s, m, i) \geq 0$. Initially, $V_0(r, s, m, i) \equiv 0$, thus $V_1(r, s, m, i) - V_0(r, s, m, i) \geq 0$. Since $\max(a_{n+1}, b_{n+1}) - \max(a_n, b_n) \geq \min(a_{n+1} - a_n, b_{n+1} - b_n)$,

$$V_{n+1}(r, s, m, i) - V_n(r, s, m, i) \geq \text{Min} \left\{ \frac{(m+m_0^i)}{(s+s_0^i)} \beta [V_n(r, s+1, m+1, i) - V_{n-1}(r, s+1, m+1, i)] + \left(1 - \frac{(m+m_0^i)}{(s+s_0^i)}\right) \beta [V_n(r, s+1, m, i) - V_{n-1}(r, s+1, m, i)], V_n(r(1-r)F_i(m), s_0^{i+1}, m_0^{i+1}, i+1) - V_{n-1}(r(1-r)F_i(m), s_0^{i+1}, m_0^{i+1}, i+1) \right\}$$

(A21)

Hence, the induction step holds and $V_n(r, s, m, i)$ is non-decreasing in n . Thus, being a non-decreasing but bounded set of functions, by Puterman (2005), it converges to $V(r, s, m, i)$. Therefore, the result is proven. Now we are in a position to prove the way the optimal value function varies as the different components of the state space vary.

Lemma 2:

- i) $V(r, s, m, i)$ is non-increasing in r ;
- ii) $(1-r)f_i(m+1) + \beta V(r, s, m+1, i) \geq \beta V(r, s, m, i)$
- iii) $V(r, s, m, i)$ is non-increasing in s (i.e. $V(r, s, m, i) \geq V(r, s+1, m, i)$);
- iv) $V(r, s, m, i)$ is convex in r .

Note that (ii) says that the recovery rate is higher if there is a payment in the current period than if there is no payment in the period.

Proof: All the proofs are by induction on the iteration number in the value iteration scheme (2).

In each case, the result is trivially true for the 0th iteration since $V_0(r, s, m, i) \equiv 0, \forall r, s, m, i$.

i) Assuming $V_n(r, s, m, i)$ is non-increasing in r , then as $r + (1-r)F_i(m)$ is increasing in r , all the terms on the RHS of (2) are non-increasing in r and hence so is $V_{n+1}(r, s, m, i)$.

ii) This holds trivially when $n = 0$ since $V_0(r, s, m, i) = 0$ for all values of r, s, m and i . To prove the induction step from n to $n+1$, we use the inequality

$$\max(a_{n+1}, b_{n+1}) - \max(a_n, b_n) \leq \max(a_{n+1} - a_n, b_{n+1} - b_n) \quad (\text{A2})$$

The induction step then becomes:

$$\begin{aligned} & V_{n+1}(r, s, m, i) - V_{n+1}(r, s, m+1, i) \\ & \leq \text{Max} \left\{ \begin{aligned} & \frac{(m+m_0^i)}{(s+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+1, m+1, i)] - \frac{(m+1+m_0^i)}{(s+s_0^i)} [(1-r)f_i(m+2) + \beta V_n(r, s+1, m+2, i)] \\ & + \left(1 - \frac{(m+m_0^i)}{(s+s_0^i)}\right) \beta V_n(r, s+1, m, i) - \left(1 - \frac{(m+1+m_0^i)}{(s+s_0^i)}\right) \beta V_n(r, s+1, m+1, i); \\ & V(r + (1-r)F(m), s_0^{i+1}, m_0^{1+1}, i+1) - V(r + (1-r)F(m+1), s_0^{i+1}, m_0^{1+1}, i+1) \end{aligned} \right. \\ & = \text{Max} \left\{ \begin{aligned} & \frac{(m+1+m_0^i)}{(s+s_0^i)} [(1-r)f_i(m+1) - (1-r)f_i(m+2) + \beta V_n(r, s+1, m+1, i) - \beta V_n(r, s+1, m+2, i)] \\ & + \left(1 - \frac{(m+1+m_0^i)}{(s+s_0^i)}\right) \beta [V_n(r, s+1, m, i) - V_n(r, s+1, m+1, i)] \frac{-1}{(s+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+1, m+1, i) - \beta V_n(r, s+1, m, i)]; \\ & V(r + (1-r)F(m), s_0^{i+1}, m_0^{1+1}, i+1) - V(r + (1-r)F(m+1), s_0^{i+1}, m_0^{1+1}, i+1) \end{aligned} \right. \\ & \leq \text{Max} \left\{ \begin{aligned} & \frac{(m+1+m_0^i)}{(s+s_0^i)} [(1-r)f_i(m+1)] + \left(1 - \frac{(m+1+m_0^i)}{(s+s_0^i)}\right) [(1-r)f_i(m+1)] \\ & \frac{-1}{(s+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+1, m+1, i) - \beta V_n(r, s+1, m, i)]; \\ & (1-r)f_i(m+1) \end{aligned} \right. \quad (\text{A3}) \\ & \leq (1-r)f_i(m+1) \end{aligned}$$

Again the first inequality follows from (A2). The second inequality follows from applying the induction hypothesis (ii) to each bracket in turn. This shows that the bracketed part of the third term is in fact positive while the first two terms sum up to $(1-r)f_i(m+1)$. The inequality in the second expression follows since $V(r, s_0^{i+1}, m_0^{1+1}, i+1) - V(r+h, s_0^{i+1}, m_0^{1+1}, i+1)$ cannot exceed h since that is the maximum recovery rate difference that is possible.

iii) We also prove the monotonic drop in the value function as s increases using induction on the n in value iteration. It is enough to show $V_n(r, s+1, m, i) - V_n(r, s, m, i) \leq 0$. Trivially this holds when $n = 0$ since $V_0(\dots)$ is zero. Assume it is true for V_n , then:

$$\begin{aligned}
& V_{n+1}(r, s+1, m, i) - V_{n+1}(r, s, m, i) \\
& \leq \text{Max} \begin{cases} \frac{(m+m_0^i)}{(s+1+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+2, m+1, i)] - \frac{(m+m_0^i)}{(s+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+1, m+1, i)] \\ + \left(1 - \frac{(m+m_0^i)}{(s+1+s_0^i)}\right) \beta V_n(r, s+2, m, i) - \left(1 - \frac{(m+m_0^i)}{(s+s_0^i)}\right) \beta V_n(r, s+1, m, i), \\ 0 \end{cases} \quad (\text{A4}) \\
& = \text{Max} \begin{cases} \frac{(m+m_0^i)}{(s+1+s_0^i)} [\beta V_n(r, s+2, m+1, i) - \beta V_n(r, s+1, m+1, i)] + \left(1 - \frac{(m+m_0^i)}{(s+1+s_0^i)}\right) \beta [V_n(r, s+2, m, i) - V_n(r, s+1, m, i)] \\ \frac{-(m+m_0^i)}{(s+1+s_0^i)(s+s_0^i)} [(1-r)f_i(m+1) + \beta V_n(r, s+1, m+1, i) - \beta V_n(r, s+2, m, i)], \\ 0 \end{cases} \\
& \leq 0
\end{aligned}$$

The first inequality follows from (A2). The second inequality follows since the first two terms in the first expression are negative because of the induction hypothesis on (iii), while the third term is negative because of the induction hypothesis of (ii) holding for $V_n(\cdot)$. This means that the bracketed part of the third term is positive. Thus, the difference $V_{n+1}(r, s+1, m, i) - V_{n+1}(r, s, m, i)$ is less than or equal to zero and $V_{n+1}(\cdot)$ is non-increasing in s . Hence, the induction assumption holds.

iv) Since $V_0(r, s, m, i) \equiv 0, \forall r, s, m, i$, the induction hypothesis that $V_n(r, s, m, i)$ is convex holds for $n=0$. Examining equation (2), the first expression on the RHS is the sum of a linear function in r and two functions which are convex from the induction hypothesis. Therefore, the expression is convex in r . Similarly, the induction hypothesis means the second expression is convex. Since the maximum of two convex functions is convex, the induction is proven and the result holds.

There is no monotone property for m , because, as this increases, one is more likely to get more repayments with this course of action, but there is less money to recover from the future courses

of actions. These two factors work against each other and negate the possibility of a monotone result.

Proof of Result 1: From Lemma 2(ii), $V(r, s, m, i)$ is non-increasing in s . Let:

$$s^*(r, m, i) = \min\{s \mid V(r, s, m, i) = V(r + (1-r)F_i(m), 0, 0, i+1)\},$$

which means the optimal policy is to move to action $i+1$. The minimum must exist since, if $V(r, s, m, i) > V(r + (1-r)F_i(m), 0, 0, i+1)$, then as s tends to infinity, $V(r, s, m, i)$ tends to 0.

This means the inequality $V(r, s, m, i) > V(r + (1-r)F_i(m), 0, 0, i+1)$ cannot continue to hold.

For $s > s^*(r, m, i)$, we have:

$$V(r + (1-r)F_i(m), 0, 0, i+1) \leq V(r, s, m, i) \leq V(r, s^*(r, m, i), m, i) = V(r + (1-r)F_i(m), 0, 0, i+1)$$

where the middle inequality follows from the non-increasing property in s . Hence, in state (r, s, m, i) , one must also choose action $i+1$.

Proof for Result 2: It is sufficient to show that, when in state (r, s, m, i) it is optimal to stay with action i , then it remains optimal to do so in state $(r, s, m+1, i)$. Let π be the optimal policy starting in state (r, s, m, i) . The total return under this policy, $V^\pi(r, s, m, i)$, is decomposed into a term describing the return under action i and one describing the recovery rate when other actions are used, namely:

$$V^\pi(r, s, m, i) = W^\pi(r, s, m, i) + \sum_{s', m'} p^\pi(s', m') V(r + (1-r)(m + m')f_i, 0, 0, i+1)$$

Here $W^\pi(r, s, m, i)$ is the total recovery using action i under policy π , while $p^\pi(s', m')$ is the probability that one moves to action $i+1$ after s' more periods which involved m' more repayments. Note $\sum_{s', m'} p(s', m') = 1$. If we apply policy π starting in state $(r, s, m+1, i)$, we have a similar decomposition:

$$\begin{aligned}
& V^\pi(r, s, m + 1, i) \\
&= W^\pi(r, s, m + 1, i) + \sum_{s', m'} p^\pi(s', m') V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1)
\end{aligned}$$

Note $W^\pi(r, s, m + 1, i) > W^\pi(r, s, m, i)$, since the first term has higher probabilities of a repayment each period compared with the second.

We need to prove that if $V(r, s, m, i) > V(r + (1 - r)mf_i, 0, 0, i + 1)$, then $V(r, s, m + 1, i) > V(r + (1 - r)(m + 1)f_i, 0, 0, i + 1)$.

$$\begin{aligned}
& V(r, s, m + 1, i) \\
&\geq W^\pi(r, s, m + 1, i) + \sum_{s', m'} p^\pi(s', m') V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1) \\
&\geq W^\pi(r, s, m, i) + \sum_{s', m'} p^\pi(s', m') V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1) \\
&= V(r, s, m, i) + \sum_{s', m'} p^\pi(s', m') (V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1) - V(r + (1 - r)(m + m')f_i, 0, 0, i + 1)) \\
&\geq V(r + (1 - r)mf_i, 0, 0, i + 1) + \sum_{s', m'} p^\pi(s', m') (V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1) - V(r + (1 - r)(m + m')f_i, 0, 0, i + 1)) \\
&\geq V(r + (1 - r)(m + 1)f_i, 0, 0, i + 1).
\end{aligned}$$

The last inequality follows from the convexity of $V(r, s, m + 1, i)$ in r , since convexity implies:

$$\begin{aligned}
& V(r + (1 - r)mf_i, 0, 0, i + 1) - V(r + (1 - r)(m + 1)f_i, 0, 0, i + 1) \\
&\geq V(r + (1 - r)(m + m')f_i, 0, 0, i + 1) - V(r + (1 - r)(m + 1 + m')f_i, 0, 0, i + 1).
\end{aligned}$$

Together with $\sum_{s', m'} p^\pi(s', m') = 1$, this gives the inequality and hence the proof is complete.

Table 1. Value functions $V(0,0,0,1)$ (i.e. the expected recovery rate) for the proposed model and two simpler alternatives.

	Expected Recovery Rate
Optimal Bayesian Model	0.1926
Bayesian Model with Myopic Policy	0.1894
Non-Bayesian Policy	0.1747

Figure 1. Optimal policies for states $(0,s,m,1)$ for all s and m . The policy for states highlighted in red is to stay in Action 1 and for those in green to move to Action 2. The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.

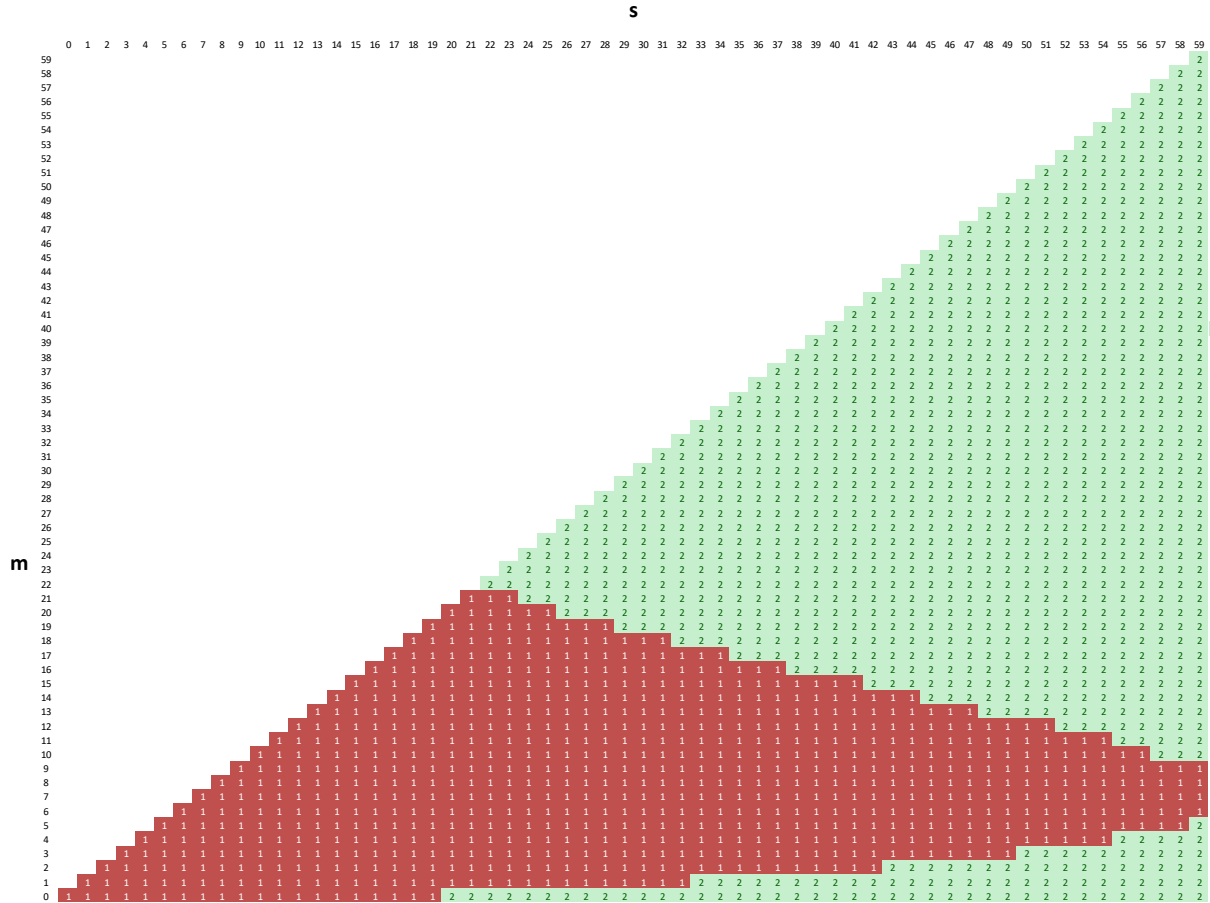


Figure 2a. Optimal policies for states $(F_1(1),s,m,2)$ (top) and $(F_1(10),s,m,2)$ (bottom) for all s and m . The policy for states highlighted in green is to stay in Action 2 and those in blue is to move to Action 3. The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.

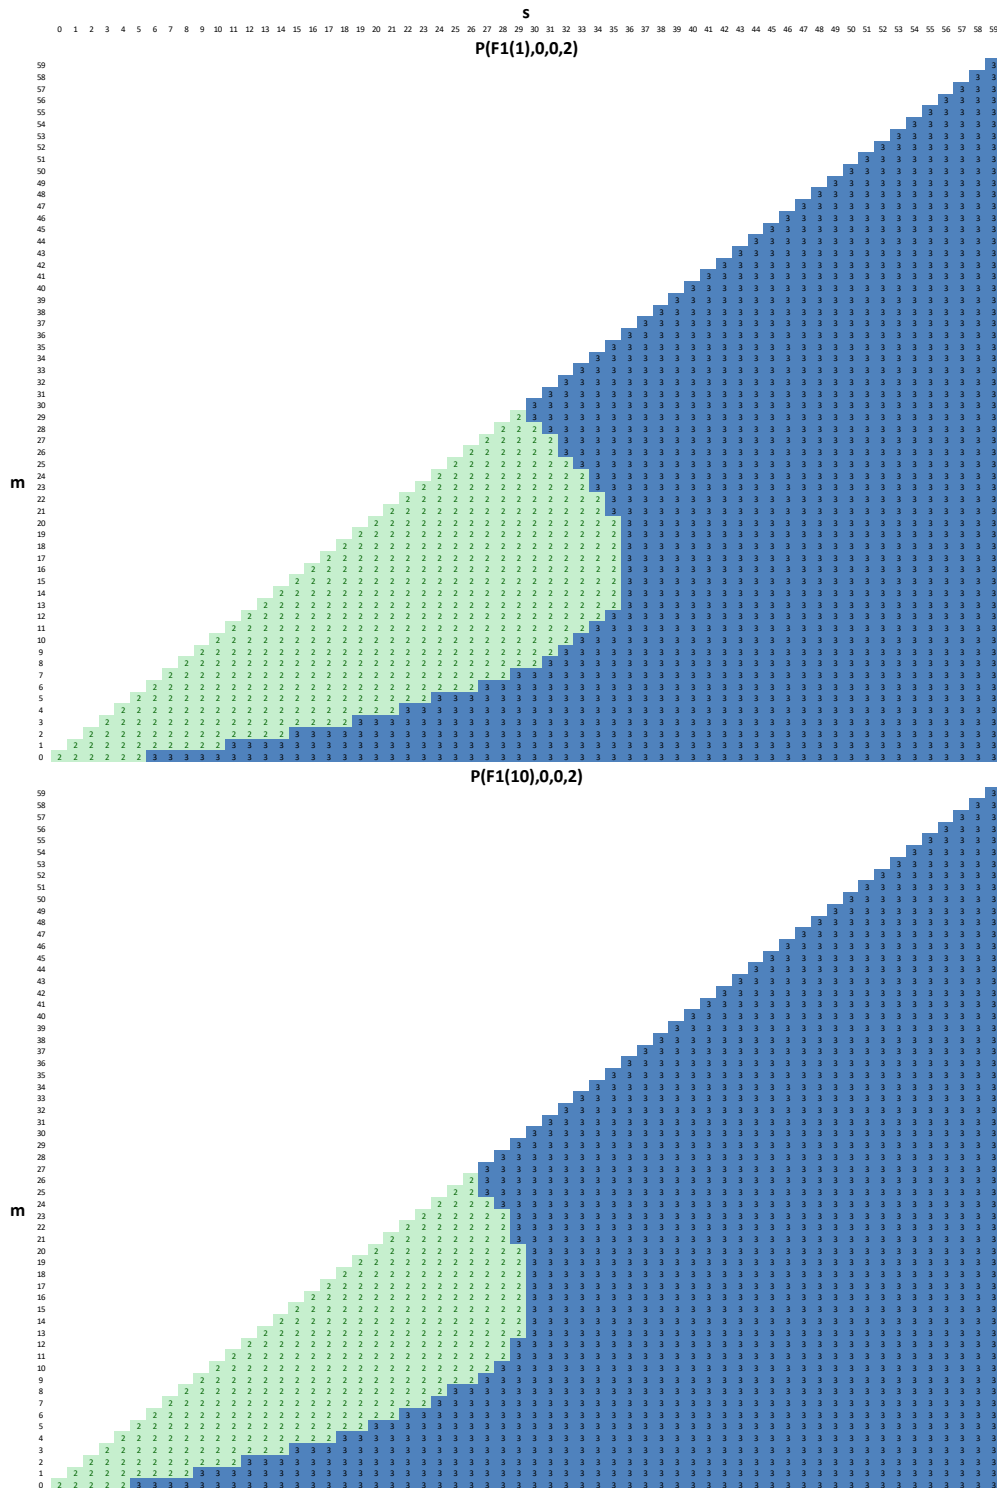


Figure 2b. Optimal policies for states $(F_1(30), s, m, 2)$ (top) and $(F_1(50), s, m, 2)$ (bottom) for all s and m . The policy for states highlighted in green is to stay in Action 2 and those in blue is to move to Action 3. The values along the axis are the values for $m(=0, 1, \dots, 60)$ and $s(=0, 1, \dots, 60)$.

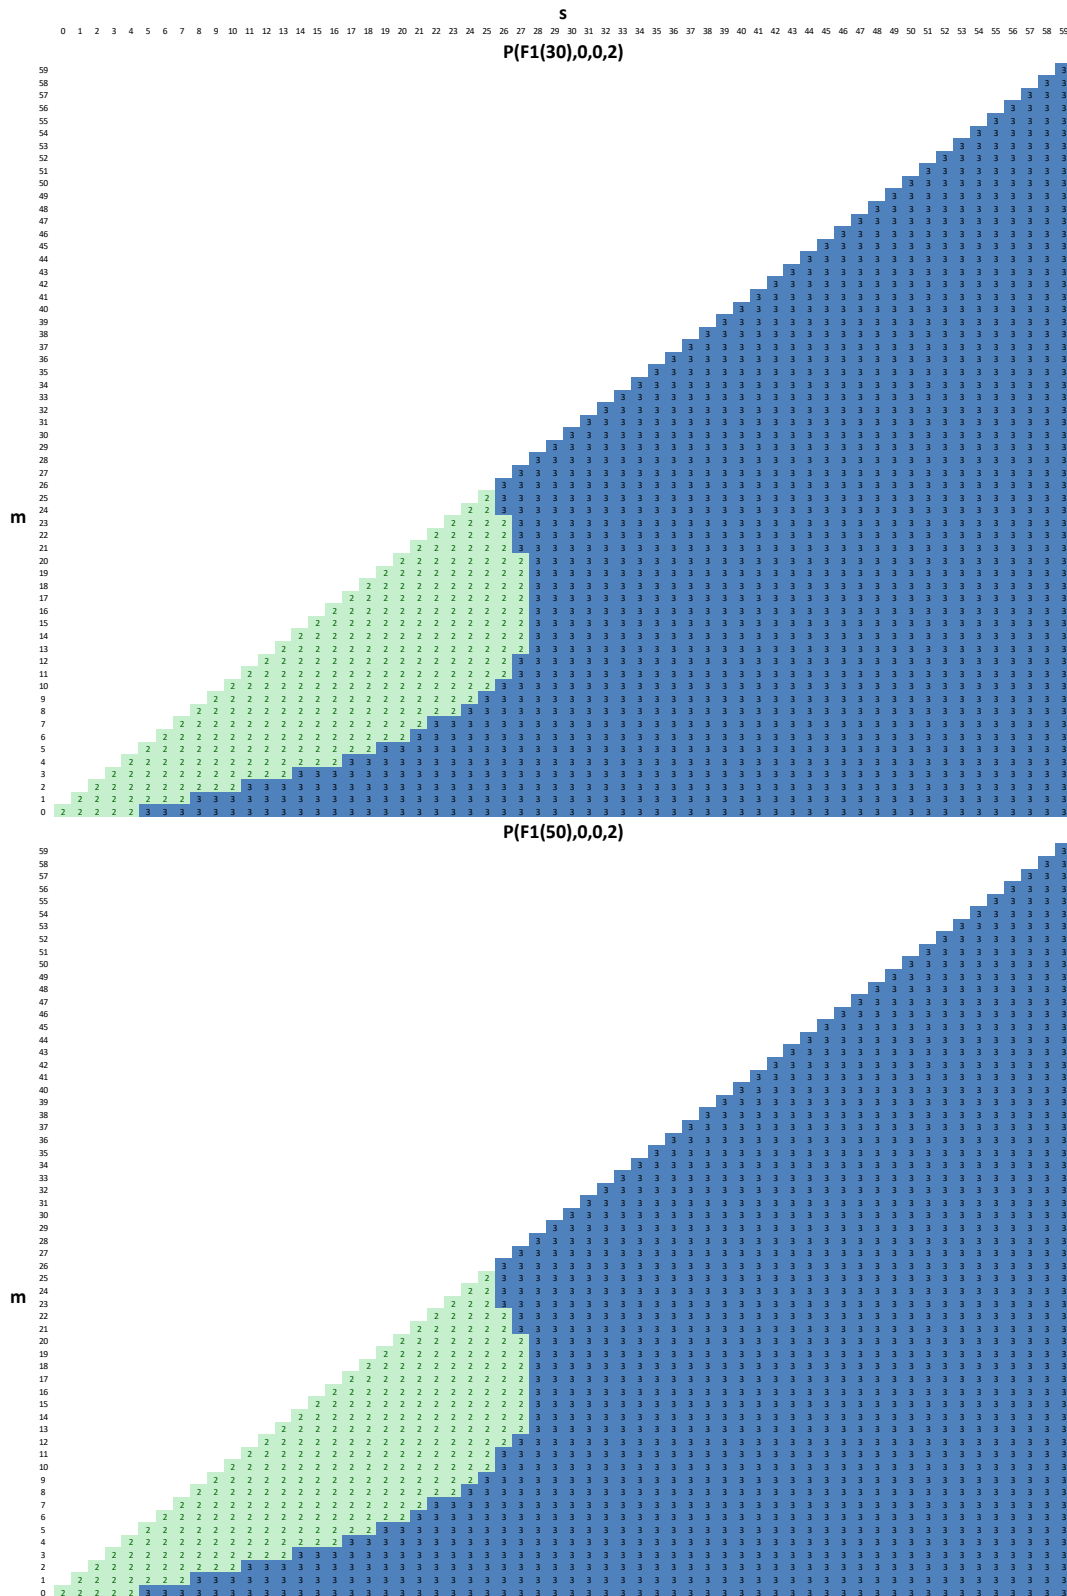


Figure 3. Optimal policies for states $(0,s,m,1)$ for all s and m , for a fixed repayment rate. The policy for states highlighted in red is to stay in Action 1 and for those in green to move to Action 2. The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.




Figure 4a. Optimal policies for states $(F_1(1),s,m,2)$ (top) and $(F_1(10),s,m,2)$ (bottom) for all s and m , for a fixed repayment rate. The policy for states highlighted in green is to stay in Action 2 and those in blue is to move to Action 3. The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.

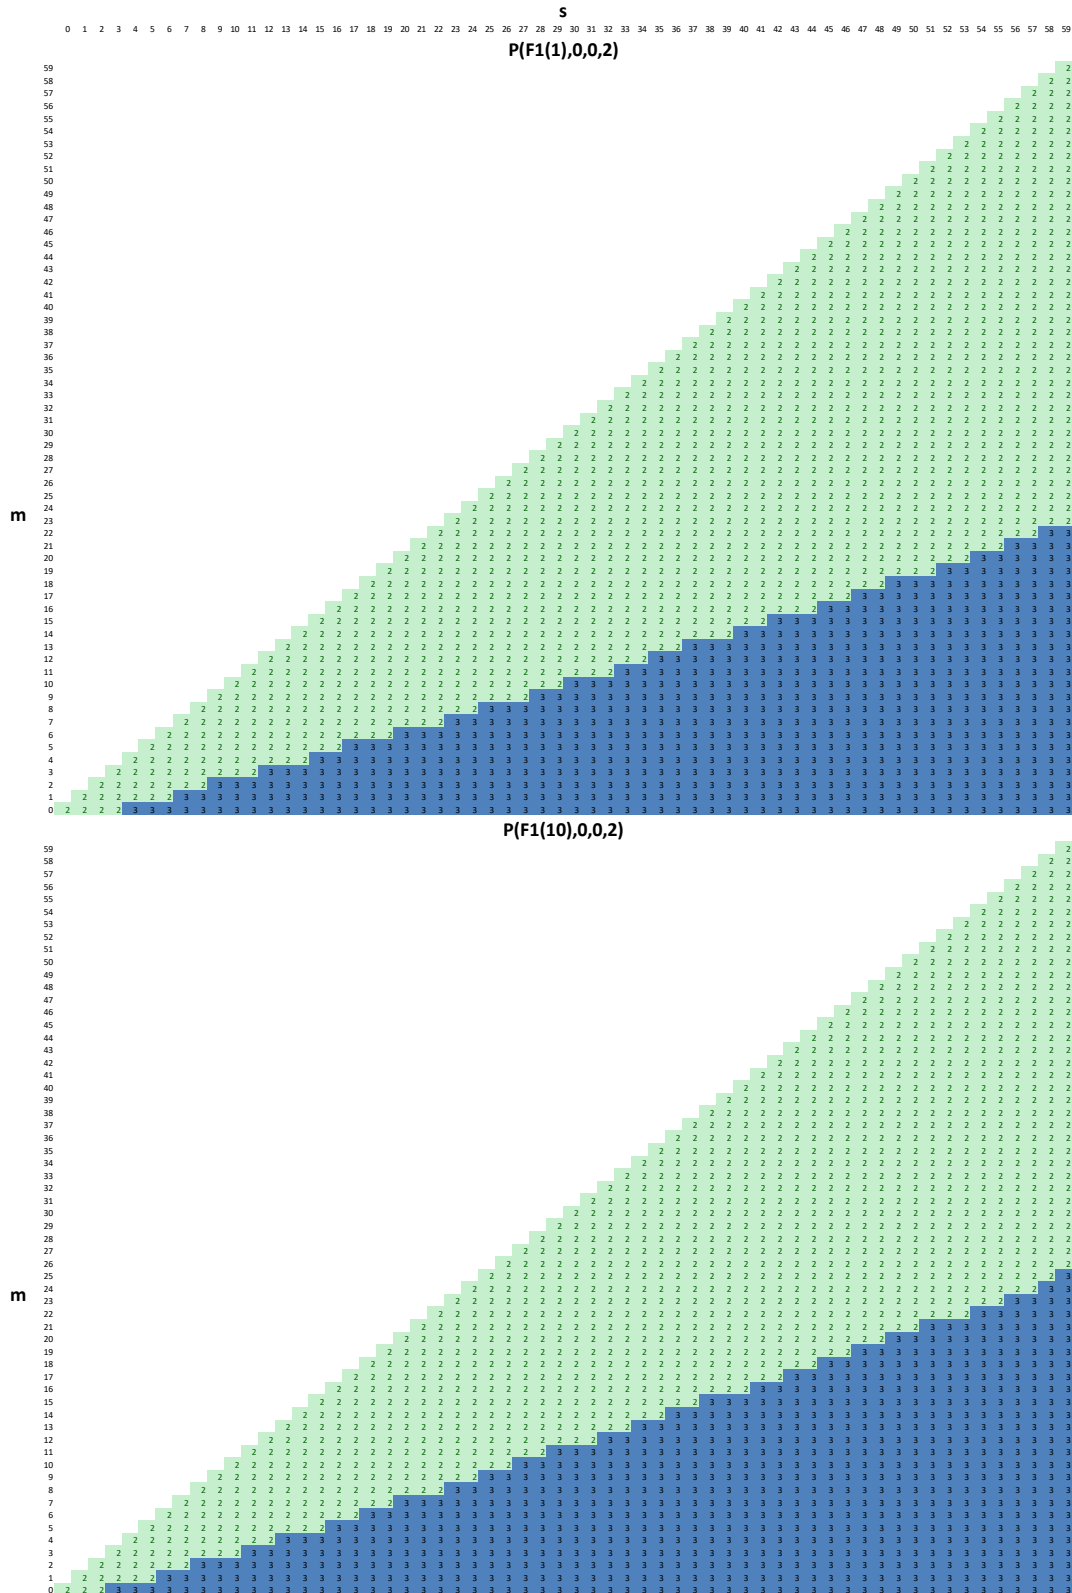


Figure 4b. Optimal policies for states $(F_1(30),s,m,2)$ (top) and $(F_1(59),s,m,2)$ (bottom) for all s and m , for a fixed repayment rate. The policy for states highlighted in green is to stay in Action 2 and those in blue is to move to Action 3. The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.




Figure 5. The myopic policies for states $(0,s,m,1)$ for all s and m . The policy for states highlighted in red is to stay in Action 1; the policy for states highlighted in green is to move to Action 2; states in blue font indicate where the myopic policy is choosing Action 1 whereas the optimal policy chooses Action 2 (as shown in Figure 1.). The values along the axis are the values for $m(=0,1,\dots,60)$ and $s(=0,1,\dots,60)$.




Figure 6a. The optimal non-Bayesian policies for states $(0,m,1)$

m	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
Policy	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2
m	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39
Policy	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
m	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59
Policy	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2

Figure 6b. The optimal non-Bayesian policies for states $(F_1(m_1), m, 2)$, for all m_1 and $m=0,1,\dots,19$. Due to page size limitations, we have not shown the results of $(F_1(m_1), m, 2)$ for $m=20,21,\dots,60$. The optimal policies for those states are using Action 3. The policy for states highlighted in green is to stay in Action 2 and those in blue is to move to Action 3.

m	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
r=F1(0)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3
r=F1(1)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3
r=F1(2)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3
r=F1(3)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3
r=F1(4)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3
r=F1(5)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3
r=F1(6)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3
r=F1(7)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3
r=F1(8)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3
r=F1(9)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3
r=F1(10)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(11)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(12)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(13)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(14)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(15)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(16)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(17)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(18)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(19)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(20)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(21)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(22)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(23)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(24)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(25)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(26)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(27)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(28)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(29)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(30)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(31)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(32)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(33)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(34)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(35)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(36)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(37)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(38)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(39)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(40)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(41)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(42)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(43)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(44)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(45)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(46)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(47)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(48)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(49)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(50)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(51)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(52)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(53)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(54)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(55)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(56)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(57)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(58)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(59)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3
r=F1(60)	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	3	3	3	3	3