UNIVERSITY OF SOUTHAMPTON

SCHOOL OF MANAGEMENT

USING ADAPTIVE LEARNING IN CREDIT SCORING TO ESTIMATE ACCEPTANCE PROBABILITY DISTRIBUTION

by

Hsin-Vonn Seow

Submitted for the degree of Doctor of Philosophy

APRIL 2006

Correction Sheet

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

SCHOOL OF MANAGEMENT

Doctor of Philosophy

USING ADAPTIVE LEARNING IN CREDIT SCORING TO ESTIMATE ACCEPTANCE PROBABILITY DISTRIBUTION

By Hsin-Vonn Seow

Credit Scoring and Behaviour Scoring are tools that are widely used in the applications of quantitative analysis in businesses. The main purpose is to assess the risk of such customers defaulting but in the commercial environment. Currently in consumer lending, there is an increasing need to assess whether the customer is most likely to accept a variant of a product. Thus, if one wishes to use the scorecard both for risk assessment and for product acceptance, one may wish to vary the product offered and the questions asked in order to improve the estimates of the probability of acceptance as a function of the features offered.

In the beginning phase, we will look at the strategies that only change one feature of the product offered, so as to improve the overall profitability of the product. This includes assessing and updating the probability of the customer is accepting the various versions of the product. The underlying model that was built was an adaptive dynamic programming model with elements of Bayes theorem to include previous actions.

The subsequent work was to look at the possibilities of adjusting the questions asked in the scorecard to improve the estimates of customer acceptance without diminishing the credit risk assessment. Applicants for credit have to provide information for the risk assessment process. In the current conditions of a saturated consumer lending market, and hence falling take rates, information like this can be used to assess the probability of a customer accepting the offer. Also, lenders do not want to make the application process too complicated, and with the growth in adaptive marketing channels like the Internet and the telephone, they can make the questions they ask depend on the previous answers. We investigated how one could develop such 'adaptive' application forms; which would assess acceptance probabilities as well as risk of default.

Finally, we look at a model which looks both at which question to ask and what offer to make when the answers to the question is known. This extends the work of the earlier dynamic programming model.

Table of Contents

List of Figures

List of Tables

DECLARATION OF AUTHORSHIP

I, Hsin-Vonn Seow, declare that the thesis entitled "Using Adaptive Learning in Credit Scoring to Estimate Acceptance Probability Distribution" and the work presented in it are my own. I confirm that:

- this work was done wholly and mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
  - (i) Seow H. V. and Thomas L. C. (to be published). Using Adaptive Learning in Credit Scoring to Estimate Take-Up Probability Distribution. *European Journal of Operational Research.*
  - (ii) Seow H. V. and Thomas L. C. (to be published). To Ask or Not to Ask, That is the Question. *European Journal of Operational Research.*

Signed: [redacted]..

Date: 28 APRIL 2006...

ACKNOWLEDGEMENTS

I thank God for the many blessings Lee and I have received that have made the stay in Southampton a memorable one. I would like to express my gratitude and many thanks to Professor Lyn Thomas for his untiring patience (with reading the thesis and also my seemingly endless request for conferences), guidance, interest and motivation. Words cannot express what you have helped me achieve these 3 and a half years and for that I have much to be grateful to you. Thanks also to Lyn, Margery and Steve for the delicious meals and the hospitality you have showered on Lee and me these past years.

I like to express gratitude to Fair Isaac Inc. and the School of Management at the University of Southampton for funding my Ph.D. studies for the past 3 and a half years.

Special thanks to Edgar, Alex, Suchi, James, Nim and Marcus for helping us realise that there is room for socialising. The dinner/board games sessions with Edgar and Alex are always fun and I will sorely miss these times. To Christophe, what can I say, thanks for putting up with my whining and "boo-hoo" ing. You made my Ph.D. life more interesting. "Dental floss" :p Thanks also to Bart for the dinners and chats that helped me get through my difficult times on the Ph.D. To members of the ESI 2004 in Ankara, especially Stefan and Belen, many thanks for being there for me! My pub quiz team (particularly Debbie, Jayne and Clare) whose laughter and positive outlook on life has left a deep impression on me.

Special mention to Julia and Howard for their kindness and understanding. You are the best landlady/lord that we ever had!

I dedicate this thesis to my family, especially my mum and dad, for their unconditional support and love. I also dedicate this to my husband, Lee who has through everything been patient, understanding and loving.

Chapter 1
Introduction

1.0 Credit

UK consumer credit has been on the rise. In 2004, it was calculated that the consumer credit in the UK was over two and a half times greater than in 1993. Total expenditure on both credit and debit cards have increased from £209 billion in 2002 to £231 billion in 2003. To give an idea of how much consumer credit grown, according to the National Statistics department, for 2003, 90 percent of women and 91 percent of men possess at least one card (see **Figure 1.1**). In addition to that, for the same year, 87 percent of the 16 to 24 year olds and the 65 and over all had plastic cards, unlike in 2001 (**Figure 1.2**) where about 31 percent of the 16-24 year olds owned plastic cards where 66 percent of the 65 and over owned these cards.

Hence, quoting Lewis (1992), consumer credit is an essential enterprise; it is very profitable considering the volume of potential customers. Financial products like credit cards, personal loans, overdrafts and mortgages are readily available to the consumers courtesy of banks and financial institutions. Most of these financial institutions use credit scoring systems to help them make decision on which applicants to accept and which to reject.

Credit scoring is used to differentiate the "good" accounts from the "bad" ones but consumer credit was a lenders' market. The decision of whether an individual was given a financial product was solely up to the lender. But this lenders' market environment has shifted to what is now a buyers' market, due to the saturation of financial products available in the market (Purang, 2005). With the changing economic conditions, people have access to much more credit and their spending has been much more dependent on credit. More consumers will "qualify" for receiving credit or other financial products and because there are many places that offer these services, they can afford to shop around for a deal that better suits them.

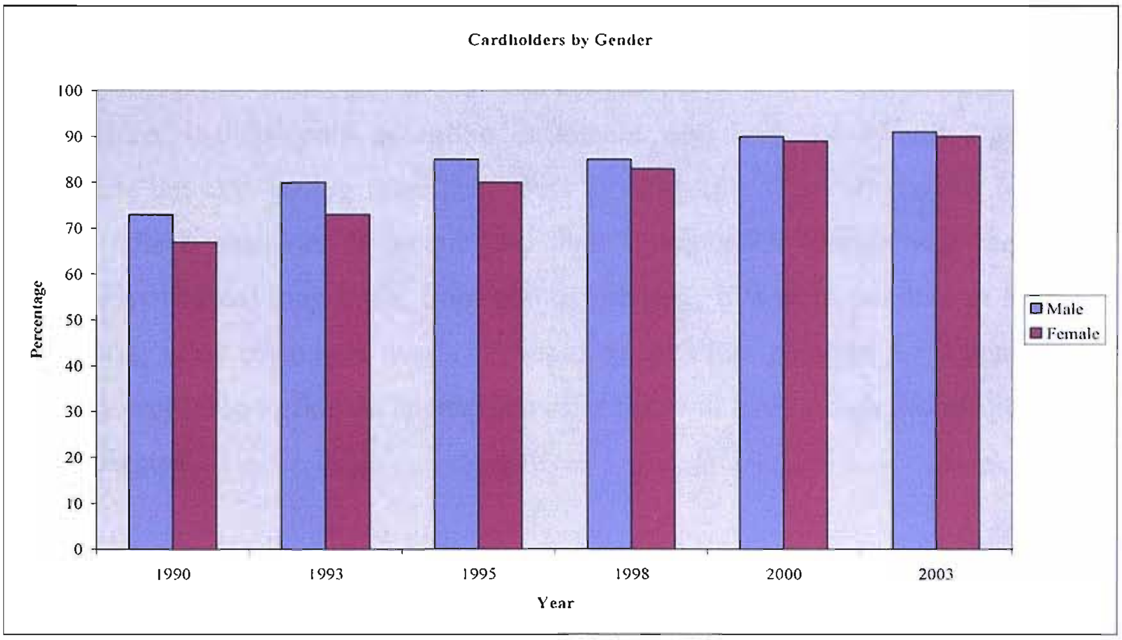

Figure 1.1: Credit card holders in the UK from 1990-2003

(http://www.statistics.gov.uk/downloads/theme_social/Social_Trends35/06_09.xls and http://www.statistics.gov.uk/StatBase/ssdataset.asp?vlnk=5007&Pos=&ColRank=1&Rank=272)

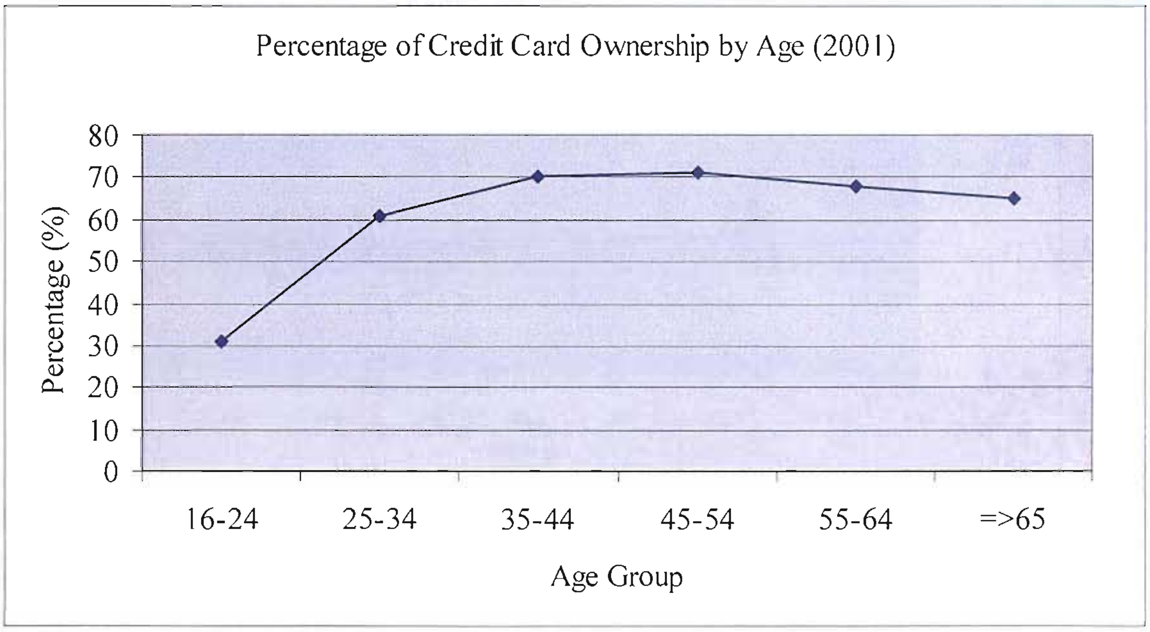

Figure 1.2: Credit card ownership by age for 2001

(http://www.statistics.gov.uk/StatBase/ssdataset.asp?vlnk=6321&Pos=1&ColRank=2&Rank=544)

Hence the rate of acceptance of offers of credit is dropping since customers now have more choices. So to continue getting profit from such a market, attention has to be placed on not only accepting customers who will not default ("good" customers) but also getting these customers to accept the financial product being offered. Each customer is unique and their living needs (which will require financial assistance) may differ from one to another. If it were possible to learn more about these customers needs, it would be therefore possible for a bank or financial institution to find an appropriate offer that will have a high probability of being accepted.

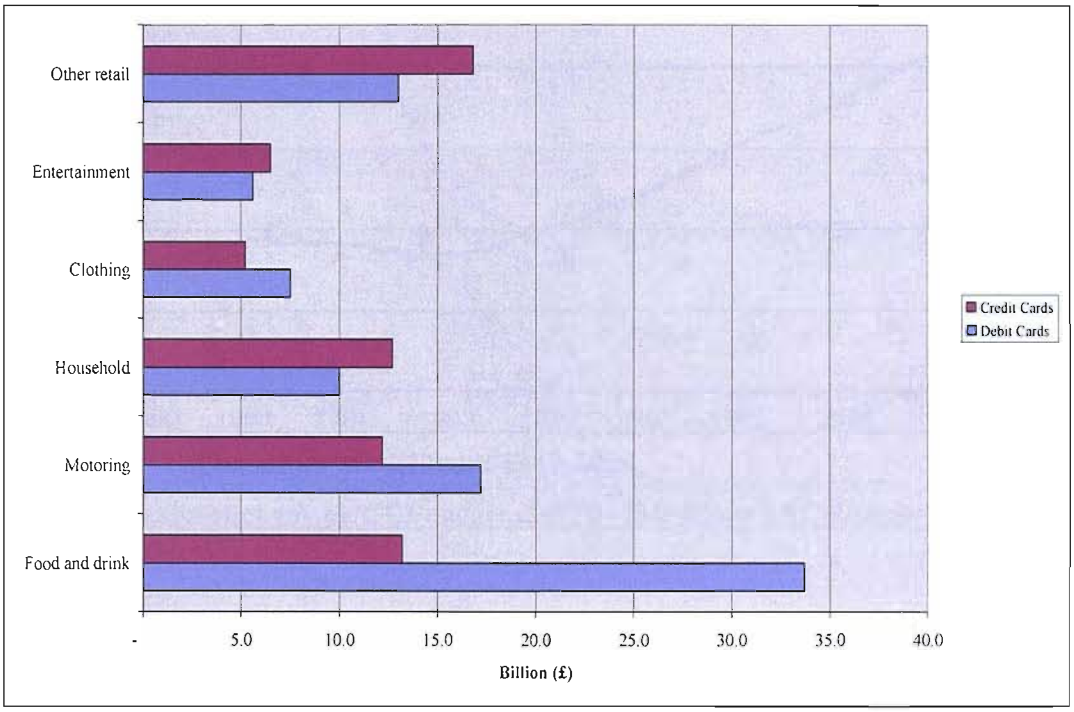

Figure 1.3: Purchase by credit and debit card for the year 2003 (http://www.statistics.gov.uk/downloads/theme_social/Social_Trends35/06_10.xls)

In order to perform any kind of customization, one must first learn about the customers, more specifically, the customers' preference. Only when we know enough, can there be a matching of offer to the customers' need. If this matching is done well, the acceptance rate of offers will increase.

1.1 Current Market Condition

It has been mentioned that the market was traditionally a lenders market. It was simply an environment where demand was greater than the supply. Hence it was the lenders that assessed whether a person applying for credit was worth offering credit to. This has since changed to an environment which we refer to as the buyers market.

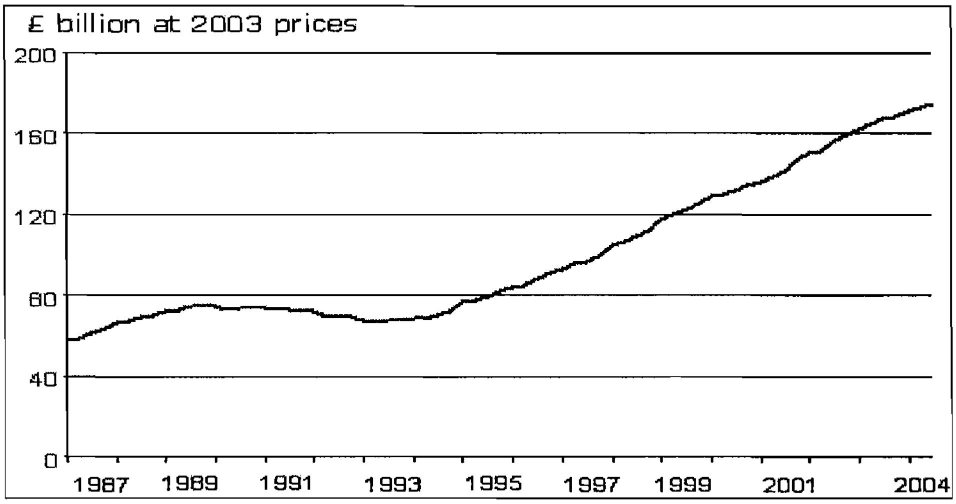

Figure 1.4: Consumer credit from 1987-2004
(http://www.statistics.gov.uk/CCI/nugget.asp?ID=1048&Pos=6&ColRank=2&Rank=224))

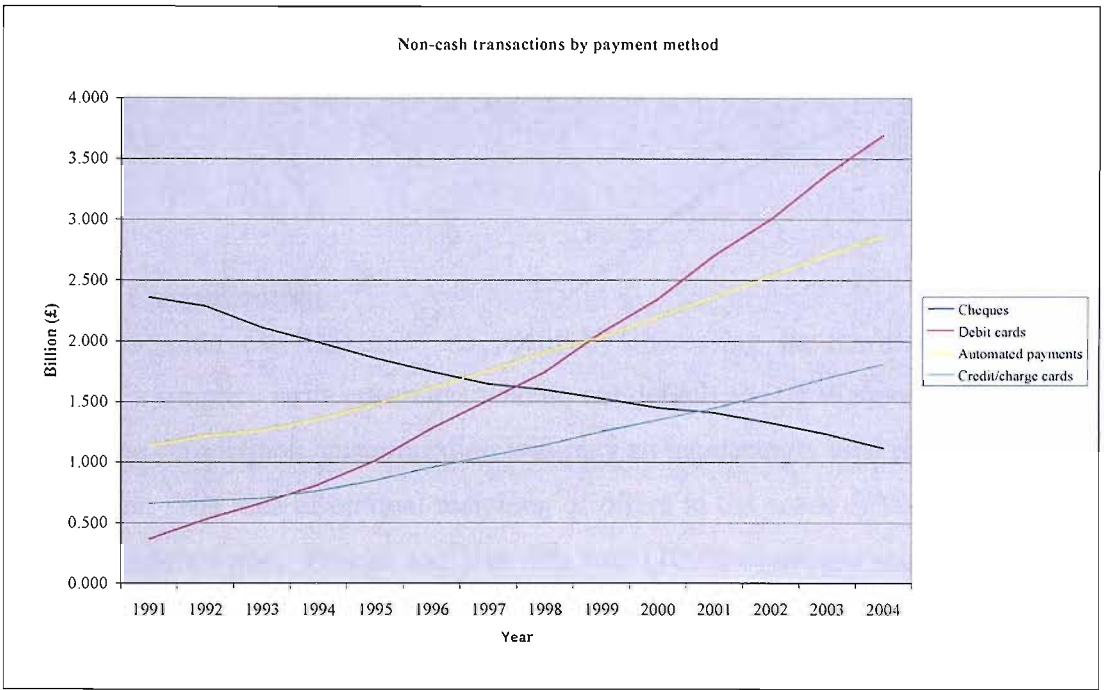

Figure 1.5: Non-cash transaction from 1991-2004
http://www.statistics.gov.uk/downloads/theme_social/Social_Trends35/06_08.xls)

Financial products are very much used as the mode of payment for basic needs, but there are now more financial institutions that offer these products. So, new tactics and strategies are deployed to catch potential customers. The birth of new marketing channels, like the internet and telephone marketing stresses the importance of real time offers. The market now is saturated with the many choices available to the customer. Hence the key to capturing the interest of customers now lies in customization. Works like those by Raghu *et al.*, (2001) and Montgomery (2001) support this. With the internet and telephone marketing, it is possible to make real-time customized offers to increase the possibility of acceptance of the offer.

In order to do so, it is important to extract enough information about the preference of the customers to match an available product to his or her needs. We agree with Murthi and Sarkar (2003) that customization is a process consisting of (a) "learning" of consumer preferences, (b) matching of offers to the customers and (c) evaluation of the "learning" and matching processes. In this thesis, we evaluate the effectiveness of the "learning" and matching processes by looking at the acceptance of the offer. Hence our objective of customization is to maximise the acceptance of offers.

1.2 Customization

An organization can maximise its profit by increasing its market share. The equation is simple – more customers who do not default add up to more profits. But with fierce competition, customization will play an increasingly vital part to capture the market. The idea of optimal matching of offers to the needs of the customer is not a far fetched one. Prinzie and Van den Poel (2006) attempted such an optimal matching method while investigating purchasing-sequence patterns for financial services by means of Markov and mixture transition distribution models with success. See also Papamichail and Papamichail (to be published). They proposed *k*-means range algorithm to allow customers to produce personalized data clusters on the products that they preferred to help enhance purchase decisions specifically for e-commerce.

But customization in this context does not mean coming up with products that precisely meets the idiosyncratic needs of each individual customer (Stump *et al.*, 2002) but more to predict and meet the future needs of a current or new customers by offering them a product from an existing or new range from the organization. To achieve this, the data on the customer's scorecard will have to be analysed to identify potential need. Updated questions that are used to get information to update the profile of the customer would give hints on changes in lifestyle, financial status that would help the organization decide whether to extend offers of products to the customers with an added advantage. The organization would know that the probability of acceptance would be high.

One might argue that precise customization will influence retention of current customers and also capture new ones. Although a perfectly customized product would be ideal, it would not be as profitable for the organization. The extra effort to customize a product for one customer may not be appreciated by another customer. The probability of a different customer accepting an offer that was particularly customized for one customer would not be encouraging. Further more, it will cost far more to customize a product for each individual customer compared with the meagre profit this exercise will generate.

Trying to customize an offer to fulfil the needs of customers within the range of products offered by the organization will also make the customer feel appreciated and gives an overall personal touch.

1.3 "Learning" as to Customize and Increase Acceptance

But before one can customize or personalize a product, one needs to identify the characteristics or preference of the customers (Montgomery, 2001). These characteristics can be identified through statistical analysis and this is also where we will learn about the customers. The first step in this process is the action of collecting the data from the customers.

There are many ways of obtaining such data. Either it is elicited directly through a survey or a registration form, or inferred based on past behaviour. After initial pre-processing of the data, statistical methods will then be used to analyse the information to help build a predictive model to estimate acceptance probabilities. We included the Bayesian updating element so that we could better estimate the acceptance probability of the customers and also allow the learning of how to encourage acceptance of the next offer.

Since the collection of data is an important precept of the exercise, much thought has to be put into the collection of the data. As it is important how the data is used to infer or collect data on the behaviour of customers, it is equally or more important how the data is collected in the first place.

The data we needed for this model had to have information on the applicant characteristic, the offer characteristic and also the final decision (accept or reject the offer extended) of the customer. To our knowledge, there is no publicly available data on the offer strategies of the personal finance sector. So for our model, we used the application form for a Fantasy Student account that is located online and available upon logging on to the University of Southampton student e-mail account. All the information mentioned earlier is located on a database used to store the information collected since the Fantasy Student account was started in 2001.

1.4 Question Selection

The key is to get up-to-date information that hint on any changes in the lifestyle of preference of the customer. The fastest and by far, the easiest way to do so is through asking the "right" questions. But customers will not respond well if bombarded by a long list of questions. Hence it is important to ask few but pertinent questions.

Given modern application channels such as the internet and the telephone, the process of acquiring information from applicant is a sequential one. A chain of questions, where each new question asked is determined by the response of the customer to the previous one. This approach is difficult if the medium used was a paper application form.

This chain of question asking starts by identifying a "guide" beginning question that will steer how the second question will go. The following question is selected based on the response of the customer to the previous question. The idea is to find a preference hierarchy to generate an offer that will have a high acceptance.

Recent work on this includes that by Holloway and White III (2003) on question selection for multi-attribute decision aiding, deal with guiding the decision maker to the best question to ask next to reach an ideal alternative.

1.4.1 Classification Methods

Classification methods are used for generally two purposes; one of which being used to examine and understand the interactions between the predictor variables and the designated target variable. The other is as an accurate classifier. These two purposes are not exclusive and more often then not in cases, the aim of using a classification technique is both to predict and to understand the data structure.

Some examples of classification techniques are logistic regression (LR), discriminant analysis (DA), linear programming (LP) and decision trees (DT). The classification techniques that will be used in this piece of research are logistic regression (LR) and decision trees (DT).

Logistic regression is a multivariate statistical technique used to predict a dichotomous dependent variable from a set of independent predictor variables (X). In short, the logistic regression predicts the probability that an event will occur. The outcome of interest is usually coded as $Y=1$ while other possible outcomes are denoted as $Y=0$. In this case, the sample mean of Y is equal to the sample proportion with the success outcome $\left(\sum_{i=1}^{n}\frac{y_i}{n}=p\right)$ where y=number of subjects in the sample population of n with the success outcome. The mean of Y can take any value from $-\infty$ to ∞, But the probability of the success outcome can only take the values within the interval of $[0,1]$. Hence we cannot use the denoted probability of success, p but rather the log of the odds of the successful outcome, $\log\left(\frac{p}{1-p}\right)$.

In other words, a logistic regression model is written as follows:

$$\log\left(\frac{p}{1-p}\right)=\alpha+\beta_1X_1+\beta_2X_2+\ldots+\beta_mX_m$$

where,

$X_1, X_2, \ldots, X_m$ are the independent predictor variables, and

$\alpha, \beta_1, \beta_2, ..., \beta_m$ are the unknown regression coefficients.

We will use the results from the LR to build a scorecard. Then we also use the results from the LR to find the acceptance probabilities of the applicants in our data. We use these acceptance probabilities to build the decision tree TAROT (Chapter 4) to help decide which offer to make to applicants.

When building the decision tree, we are trying to perform question selection by using the decision tree to determine which i number of questions (i = maximum questions) are most effective in extracting information on the customer. Then, based on this information and past purchasing behaviour data, attempt to come up with an offer which has a high probability of acceptance with the customer in question.

1.4.2 Decision Trees

Decision trees are a form of classification and regression trees (CART). Although primarily used for classification purposes, it can also be used to form decision trees. CART methodology allows both categorical and continuous variables to be considered as explanatory variables. It begins splitting by considering every value of a continuous variable and every split created by every possible combination of categories for every categorical variable (Haughton and Oulabi, 1997). Then the methodology chooses one split that separates the data to "two most homogenous" parts.

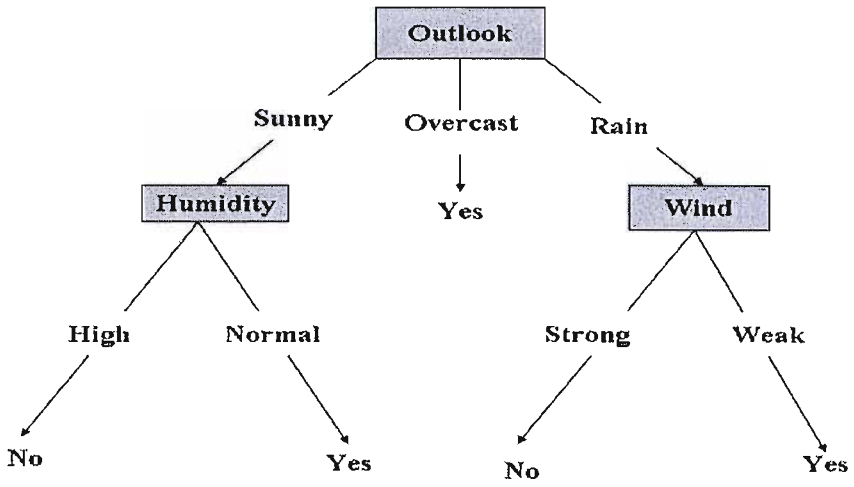

Figure 1.6: Example of a decision tree taken from T.M. Mitchells' Machine Learning (1997)

When using decision trees, there are three decisions that must be considered:

(1) Splitting rule (rule to split the set into two)

(2) Stopping rule (how to decide if a set is a terminal node)

(3) End node classification rule (how to assign terminal nodes into good and bad categories)

We will build the TAROT decision tree where we use entropy as a basis for the splitting rule. Hence the acceptance and the rejects of the offers will be perfectly divided at all levels in the tree. The height of the tree corresponds to the number of questions that want to be asked. If we want to ask two questions, then we will stop building the tree at the third level. The offer extended is the offer node with the highest acceptance probability. When using decision trees, there is always the problem of over fitting. Breiman used a method called Random Forest to overcome this problem (Breiman, 2001). This particular technique of growing individual decision trees to decide on the class by choosing the most frequently reoccurring one of the individual trees. This technique is accurate and it handles a large data set. In view of our data set from the Fantasy Student account, we try to overcome over fitting by bootstrapping. We increase the volume of our data by three-fold so that the size of the sample at a particular node is not too small.

Decision trees have been used to make decisions in the field of medicine and direct marketing. From these cases by Harper and Winslett (2006) and Thrasher (1991), it can be deduced that certain characteristics do lead to a desired action being taken. It is also believed that if situation with almost the same characteristics occurs, that past experience will be able to help decide on an appropriate course of action to achieve the desired outcome.

Hence, we will use decision trees to find how to best classify our population, and from the variables that are chosen for the splits, decide which questions to ask so that an offer which best suits the responses to the questions can be made to that customer.

In this thesis, one of the objectives is to look into the "learning" of customers' preference with an aim to aid matching of offers of high acceptance. In order to learn about the customers, the use of a minimum number of questions is examined, to "learn" enough for an offer is used to gather the essential information of customers' preference. Also, a model to calculate the probability of acceptance through dynamic programming utilising Bayesian updating is also built. The overall objective is then to connect these two tools to predict acceptance of different variants of an offer.

This model is based on adaptive decisions and dynamic programming. The questions selection uses historical data which is carefully analysed to give an offer strategy. In light of the fast moving markets, new procedures involving the need to be able to adapt quickly and automatically will be crucial. Hence there is a need to feed in information of whether previous applicants of a financial product accepted or rejected an offer into the decision process. With knowledge of the sequence of decisions, one is able to adapt the offer strategy accordingly to response, aided by adaptive dynamic programming.

This thesis is divided into the chapters as follows: in Chapter 2, we will look at customization and learning. The acceptance model complete with experimental results, is discussed in Chapter 3 while in Chapter 4, we look at the process of building TAROT using the Fantasy Student account data. In Chapter 5, we extend the acceptance model of Chapter 3, by investigating which questions to ask during the interview as well as which offer to make. Thus, the TAROT approach of Chapter 4 is a way of using past information on questions asked and offers made to identify the strategy of which question to ask and which offers to make for a large number of future applicants. So the strategy is only occasionally updated. Chapter 5, on the other hand is an introduction to a system where the strategy of which question to ask and what offers to make is updated often after a new applicant. Although this is more difficult to implement, it can respond faster to changes in the market environment more quickly than the TAROT approach. The conclusions can be found in Chapter 6.

The main contributions of this thesis are the approaches to acceptance scoring with emphasis on learning about the offers and customers so as to match the right offer to the right customer. We first built an acceptance model to predict which offer to extend to the next customer. The acceptance model is an adaptive model based on dynamic programming with Bayesian updating to include the effect of past actions in the optimality conditions (see Chapter 3). We then developed a decision tree (TAROT) to help decide which offer to make to customers with certain applicant characteristics by learning from asking questions (see Chapter 4). Finally, we extended the acceptance model to decide on the best offer to extend to the next customer by helping select a question and deciding the best offer to match each response to the question selected (see Chapter 5).

Chapter 2
Literature Review

2.0 Introduction

2.0.1 Credit Scoring

Traditionally, the lenders assess the customers' characteristics to make a decision on whether to give credit or not to someone. Before credit scoring was introduced, lenders depended on a vague and subjective rule of thumb governed by the three Cs: Character, Collateral and Capacity. This subjective method of decision making was replaced by a more reliable, statistically-based technique known as credit scoring. Lewis (1992) dates credit scoring back to the 1950s where it was first used in mail order firms as a method of finding potentially profitable customers with an aim to cut back on marketing costs. Its reliability has proven its worth as the credit scoring industry has since grown into a very successful area of consumer credit.

Credit scoring utilises statistical tests to differentiate groups in a population. It is not controlled by emotions or family ties and acquaintances, but rather on the reliable tests that will separate for example, the good accounts from the bad ones using the statistics. It is a system that is less biased than the three Cs that are based on human judgement, as the sole deciding factor. Initial results from using this technique were very encouraging as Myers and Forgy (1963) reported. Then, when the Equal Credit Opportunity Acts (ECOA) was passed in the US in 1975 and 1976, credit scoring became a vital component in decision making especially at banks and financial institutions. This meant the ban on any discrimination during the granting of credit unless that discrimination was "derived empirically and found to be statistically valid" (Thomas *et al.*, 2002).

The success of the credit cards encouraged banks to use credit scoring on their other products like personal loans in the 1980s and then later home loans and small business loans (Thomas, 2000). It is now a vital risk management tool used widely in finance and banking. Although credit scoring is used widely in industry, but most of the time, the detailed models are rarely published since with the sheer volume of lending, a small improvement in performance can lead to a substantial increase of profit (Crook *et al.*, 2001).

In the last decade, the idea of using scoring to target direct marketing campaigns has also become very popular – 50 years after it was first introduced by Sears (Lewis, 1992). This exhibits the robustness of the technique of credit scoring.

2.0.2 Techniques Used in Credit Scoring

Credit scoring involves many quantitative techniques to establish a decision on credit for a potential customer. Logistic regression (see Ohlson, 1980) and linear programming were introduced in the 1980s and are still some of the most commonly used techniques in credit scoring. Discriminant analysis, linear regression and logistic regressions are amongst the other techniques in credit scoring. Technological advances specifically ones in computing have increased the number of techniques one can use in credit scoring thus improving it. Neural networks (NN) and other expert systems are some examples of the new generation of techniques used in credit scoring. Recent works by Baesens *et al.* (2003) illustrate how support vector machines are used in distinguishing the bad and good accounts.

Other than using the predictive techniques mentioned in the previous paragraph, we will also be using decision trees (DT) to help identify the characteristics or variables that are predictive of the acceptance of an offer. We run a Classification and Regression Tree (CART) analysis to build a decision tree to help decide on an offer to extend. We take the splits from the tree and formulate questions based on the characteristic to ask the customer. In order for this to work, we limit the responses to the questions to binary responses so to correspond to the splits of the variable in the decision tree.

The tree is modelled to split the variables using entropy to ensure the best splits for the 25 applicant and offer characteristics or variables in the Fantasy Student

Account data set. The objective of the tree is to find the offer that has the highest probability of acceptance. The structure of the decision tree allows us to focus the search for the offer by classifying the data according to the applicant characteristics first, followed by a search of the best offer for that classification.

There are many offer characteristics that can be varied to capture the attention of a customer. The lender could vary the interest rate, credit limit, interest discount, insurance protection, air miles, annual fees and etc. All these different offer characteristics will appeal to different people. Some might find one of more offers tempting and have no interest in the others. Hence there is a need to classify the data set accordingly so as the "match" the most suitable offer to the appropriate set of customers.

The lender could vary these offers according to the customers to suit their financial needs and financial situation. But in order to do that, the lender must first learn about the customers. The lender will have to go back to the data collected on the customers to learn about their current financial situation to find out what kind of offers they are likely to accept at this point of time. The data is likely to have been collected during the application process either on paper, or even through online applications on the internet. What we propose is to run a logistic regression on the data set and build a scorecard for this data set. From the scores, we can obtain an acceptance probability distribution. We can also obtain the acceptance probabilities for a certain type of customer towards different variants of an offer characteristic, and so better decisions of what type of offer to extend next.

The objective of credit scoring is always to try to minimise the chance a customer will default. However, there has been a shift in the objective to how to maximise the profit made from a customer. Hence, there is a growing demand for systems and models that estimate not just the risk of default for a customer, but also how much profit the customers can bring in to the lenders. Bearing in mind for that to happen, the customers will have to have accepted offers from the lenders. In order to achieve this, lenders have to study their customers' behaviour and also learn more about them. This is done so that it is possible to better match the customers' current needs to the "right" kind of financial product. This will in turn also insure the maximum amount of profit to the lenders.

2.0.3 Acceptance Scoring

In their work, Rossi *et al.*,(1996) used logistic regression to look at the influence of demographic variables as well as using both information on just one purchase and the history of purchase. They concluded that using purchasing history to guide direct marketing decisions is better that using just one purchasing observation. Heilman *et al.*(2003) in their work to determine the appropriate amount of data for classifying consumers for direct marketing purposes also came to the same conclusion. Knott *et al.*, (2002) also used history of purchase in their next to buy models for cross-selling applications which attempted to predict the product the customer is most likely to buy next. The literatures on choice marketing models use logistic regression and classification procedures to predict what the customer would like to purchase. In these cases, it is the customer who has a wide range of offers and shows a preference for the individual offers in that range. We, on the other hand, are looking at how the retailer decides on an offer which a customer will accept.

The problems in estimating the acceptance probabilities of variants of a product come from two developments in retailing. The first is the development of new channels, where customers apply for or are offered products. The internet and direct telephone marketing mean that the offer and acceptance process is more interactive and more personal than in the past and this leads to new applications of Operational Research (OR) (Geoffrion and Krishnan (2003), Montgomery (2001), Häubl and Trifts (2000)). Secondly, the drive towards customization and personalization means that retailers are developing generic products which can have many different variants.

These are several ways in which this customization process is operationalised as the survey by Murthi and Sarkar (2003) indicated. They recall the taxonomy of price discrimination in that first degree price discrimination is the ability to charge different prices for the same product to different customers, while second degree price discrimination is offering different variants of a product at different prices to different customers. It is this second degree price discrimination which we are using here. Rossi *et al.*, (1996) were also interested in second degree price discrimination but looked at the consumer's previous purchase history to target the distribution of coupons which discounted the price of the item.

We, on the other hand, are interested in which variant of the product to offer the next customer, given the acceptance record of all the customers so far, so as to both learn the acceptance probabilities and to maximise the profitability for the organisation. The learning means we want to automatically and dynamically update the consumers' acceptance probabilities.

This work was motivated by the development of using the internet to apply for credit cards, personal loans and other forms of consumer credit. In these applications there are different variants of a product that can be offered. In credit cards, one can vary the interest rate charged, the credit limit available, the annual fee, the type and period of any initial discount, and whether features like free air miles are offered. The internet or telephone application process means that other customers do not see what is being offered and the offer can be calculated in real time as a function of the customer's characteristics, given as part of the credit scoring check on risk (Lewis (1992), Thomas *et al.*,(2002)) and the information the organisation currently has on the acceptance rates. This information can be updated immediately the offer is accepted or rejected. Hence a part of this research deals with estimating the acceptance probability from past data; see Chapter 4.

There is a substantial literature on learning so as to choose the optimal overall strategy when faced making decisions over and over again on which of a number of alternatives to undertake. When the information gained on making an action refers only to the uncertainties about that action the problem is classified as a bandit problem after the two armed gambling machines of that name. The model and the idea of an index solution being the optimal strategy were developed by Gittins and Jones (1974) and there have been a number of extensions – see for example restless bandits Whittle (1998) and a number of new applications (Gittins (1979), Gittins (1989)). However the problems here are not quite bandit problems in that accepting (or rejecting) one type of offer gives information not just about that offer but also about more (less) attractive offers. Meyer and Shi (1995) performed an experiment to study learning from feedback and found in reality there was a tendency to under-experiment with good offers and over-experiment with less promising offers. Our model seeks to use adaptive learning to select offers that would maximise the expected profit over the whole set of offers made.

2.0.4 Customization

In order to allow any customization, one must first obtain information about the characteristics of the customers as this is an important first step of customization. When collecting data, one would want data that will help predict an acceptance for an offer. This can only happen if we know which data "tell" us about how a customer with the characteristic in question reacts to a variant of an offer. So, it is equally important to know what kind of data to collect (see Piramuthu (2004) and Crone *et al.*, (to be published)).

Some common methods of collecting such information are using application forms, surveys, or make an observation of customers past behaviour and actions if available. Observations on the way customers rate how certain offer characteristics are important to them reveal what kind of offers appeal to them. Murthi and Sarkar (2003) state that responses from customers can be analysed to associate the attributes of the offers with an importance weight so show how important an offer attribute is to them.

They add that some may reveal these preferences by ranking the products or choosing from a set. Hence, in a sense, one can also say that the customers also reveal their preference via ranking their preference of attributes of an offer. Hence customization can be done by identifying the preferences, and matching it on an offer. Some examples of work on personalization or customization of bundles are discussed in Montgomery (2001), Raghu *et al.* (2001), Van den Poel and Buckinx (2005) and Peltier and Schribowsky (1997).

We believe, like Thomas *et al.* (2006) that two significant changes in the consumer lending process have lead to the estimation of the acceptance probability. The first of the two is the ability or the flexibility of the lender when it comes to the financial product and the second being the development of using interactive channels such as the internet and the telephone to make offers to customers. We take the example as in Bierman and Hausman (1970). They attempted to formulate an optimal credit granting policy using a Bayesian approach to revise probability of collection as the experience is gained. They then used dynamic programming to help with deciding the policy and tried to adapt the dynamic programming to include a decision on how much credit to offer. Another example can be found in the work of Choi *et al.,* (2003). They also used dynamic programming with Bayesian information updating to obtain an optimal two-stage ordering policy.

We are trying to match, based on previous behaviour, which offer has the highest acceptance probability to the corresponding segment. The emphasis is too find the variable that most influences the target response, which is accepting an offer.

Once the required data is acquired, one must choose the most suitable and effective method or methods to analyse the data. In this case, the first step will be to identify the significant offer attributes. Now there are numerous classification algorithms that can be used to identify the significant attributes that influence the acceptance of an offer. Past literature, when comparing methods, always compared the featured method with linear regression, non-linear programming, linear discriminant analysis

(LDA) and logistic regression (LR) to name a few. Some examples include work done by Desai *et al.* (1996), Baesens *et al.* (2002) and He *et al.,* (2005).

But as logistic regression is a better technique when it comes to predicting dichotomous outcomes, this is used to identify the significant variables that influence the acceptance of offers here. Now, logistic regression does assume variation homogeneity but in the early stages of the research, the population is considered homogenous when it comes to predicting acceptance. There are many other different types of techniques that could be utilised to manage the data collected. Crone *et al.*,(2004) advocates the use of neural networks and support vector machines in dealing with the matching of customer behaviour and preferences and product (Crone *et al.*, 2004). Baesens *et al.* (2003) uses support vector machines in the classification of good and bad account for the purpose of credit scoring with success (Baesens *et al.*, 2003).

The purpose of applying all these techniques is to find the variables that influence the acceptance of the offers. All the techniques mentioned before do perform this task but with varying accuracy. But because it is a dichotomous outcome, that is accept or reject, logistic regression is a more preferred choice.

But that is not the objective of this piece of research. By identifying the significant variables that influence the acceptance of offers, a decision tree will be utilized to identify the variables which strongly influences an accept of offer outcome. One of the advantages of the decision tree is that it is able to deal with the interactions between the variables and aid decision making by showing the probability of success for all the alternatives. Decision trees can be as powerful as most of the classification techniques like linear discriminant analysis (LDA) and logistic regression (LR). When Lee *et al.* (2006) used classification and regression trees (CART) and multivariate adaptive regression splines (MARS) to mine customer credit, they found that regression trees can outperform traditional discriminant analysis, logistic regression, neural networks and support vector machines (Lee *et al.*, 2006).

With the identification of such variables, questions with a yes/no response can be formed. The path that has the strongest probability of success will aid question selection. This means that the significant data required to process an "accept" can be acquired by only asking two or three questions, rather than 20 or more. Thus, it would be a question of asking questions corresponding to the relevant independent variable, thus minimizing the number of questions that will be asked.

This will not only be more efficient for the organization but it will not burden the customers or make them feel uncomfortable about the questions. An example by Raghu *et al.* (2001) did attempt this via influencing the information acquisition process using the questionnaire design. This was also evident in the work of Holloway and White III (2003). The difference is the usage of decision trees in building the questionnaire which will be more effective. Hence, in short, the task of acquiring the data is one consisting of sequential decisions.

2.1 Sequential Decisions

Sequential decision process describes an activity that consists of a sequence of actions taken to achieve a goal, often under uncertainty (Denardo, 1982). Thus many activities we encounter daily consist of a sequence of decisions. Examples range from driving to work to formulating credit policy models (Mehta, 1968). Customization requires a series of responses, matching of offers before a decision on an offer is made. Thus, customization is a sequential decision process. The whole process of customization is aimed at increasing the acceptance probability of an offer. It has been said that in order to achieve customization, information has to be collected on the preferences to classify potential customers into clusters. There are many ways of acquiring such information; directly as through application forms and questions, or indirectly through inference. But using either, the "next" question is always influenced by the response to the current question. This pattern of information collecting is repeated until enough information has been acquired.

The goal here is to encourage acceptance of an offer where the probability of acceptance is unknown. To achieve this goal, information about the customer's preferences has to be acquired. In the past, application forms or inferring from available information were used to acquire this information (Raghu *et al.*, 2001). We can use past information to help classify the customers and use the internet and telephone marketing, to acquire the information needed in real time and match the responses to the questions asked with an offer. To keep an edge on the competition, it is important to be able to match an offer with high acceptance probability by asking as few questions as possible.

When analysing a sequential decision process, certain characteristics do recur quite often. Hence, a collection of mathematical tools, known as dynamic programming, is used to effectively analyse the decision process. Dynamic programming's under laying principle is known as Bellman's principle of optimality (Bellman, 1957). This will prove useful for the explanation of the results in Chapter 3 and Chapter 5.

Bellman's principle of optimality states,
"An optimal policy has the property that whatever the initial state and the initial decisions are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision."

Therefore the sequential decision program can be formulated into an equation which has to be solved to obtain the value of the optimal solution. The optimal values will help find the optimal actions to take in different states of the problem. Dynamic programming is robust and can be used in many different problems as shown in Thomas (1994) and Gönül and Shi (1998).

The acceptance model in Chapter 3 is built using dynamic programming with Bayesian updating to help it "learn". We use dynamic programming to identify the optimal offer strategy while the usage of Bayes theorem here allows previous responses to be included in the decision process. It is believed that by including previous responses, the model will perform better (Van den Poel and Buckinx, 2005). Hence the popularity of methods with memory, such as neural networks (NN) and support vector machines (SVM)(see Lee *et al.*, 2006, Baesens *et al.*, 2003, Crone *et al.*, 2004 and Van Gestel *et al.,* to be published).

It is also very important to update the value of the optimal solution as the model selects the optimal action for the next state by calculating this value, which in this research is the expected future profit of the customer stream. With this updating, there will be continuous "learning" of the customers' actions, which will result in high acceptance behaviour.

Another method considered in developing stochastic methods is based on the multiarmed bandit problem. When we were working on the acceptance model at its early stages, it was thought that this lender's decision problem was very much like a two-armed and potentially multiarmed as well as the restless bandit problem made famous by Gittins (1989) and Whittle (1998).

Gittins (1989) successfully established indices that solve multiarmed bandit problems while Whittle (1998) extended these indices for the restless bandit problem. Initially, the acceptance model did look very much like a bandit problem. Two offers were tested as a two-armed bandit problem; one arm representing a Variant 1 on a credit card, and the other arm representing Variant 2 on the same credit card.

The situation seemed ideal, starting off with the Variant 1 arm and not knowing which arm to pull next. Pulling the same arm allowed the lender to learn about that particular arm; that Variant 1 was being accepted by potential customers. This "learning" allowed the lender to know when to switch the offers to Variant 2.

The probability of events happening in bandit problems is independent for the two arms. Here, the probability for one arm has to be less than the probability for the other. This is the condition applied to the two variants of offers, where the probability of accept decreases with each better variant of the offer, the expected profit obtained increases with each better variant of the offer. The better variant of the offer is the variant that will result in higher profit to the lender but has a lower acceptance probability with the customer. For example a credit card with a 10 percent interest rate on it is a better variant compared to a 5 percent interest rate charged on a credit card. It is also our belief that is a variant of the offer is accepted, a worse variant of that offer would have been accepted as well.

But unlike the bandit problems, there is a dependency between the arms. One of the assumptions of the acceptance model which will be discussed in Chapter 3 is that for a better variant of the offer (using the interest rate on credit cards example, 10 percent) to be extended, we believe that the lesser variant (5 percent interest rate) of the offer that is available would be accepted. The evolving arms in the restless bandit problem seem ideal to describe the learning from trying the two variants, but the dependency of the arms ruled out these two models.

2.2 Question Selection Method

Question-response process is also a sequential decision-making process (Holloway and White III, 2003). The data obtained from this process is used as updates and for the matching of offers. But before any updating or matching of offers can be made, it is an important precept to first acquire such data. Murthi and Sarkar (2003) listed direct and indirect ways of acquiring the data. But with today's competitive environment and the usage of the internet, offers can be made now in real time. So in order to use this to one's advantage, the choice and the number of questions one asks will play a vital role.

The data offers a range of variables that might or might not influence the acceptance of an offer. Typically, one would use LDA, LR, DT or SVM to identify the most significant variables. In comparison, the LR, DT and SVM would perform better, but the SVM takes a long time to train and between the LR and DT, the DT aids in the decision process by showing the success probabilities. This is done through a success probability attached to each significant variable, it will guide the selection of the significant variables and this will allow the formation of the needed questions. Moreover sequential partitioning procedures give slightly better classification results (see Srinivasan and Kim (1987)). Furthermore, DTs are a simple way of representing knowledge (Sörensen and Janssens, 2003). Hence, for this research, DT will be used to select the questions to be asked.

There are other methods used in making question selection. Holloway and White III (2003) tried to find a way to aid the facilitator to select the best question to ask next. They successfully developed a dynamic programming based approach based on expected optimal cost-to-go function to do so. Raghu *et al.* (2001) determined that the question asked depended on how much the lender wanted to know about the customer's preference. If it was partial information, they managed to illustrate when to stop asking questions and proved that at the half way point, the lender was better off getting full information about the customer profit wise.

The usage of DT would be just as good as other methods (Lee *et al.*, 2006) and the benefit of using the success probabilities will lead to easier interpretation of the data itself. This, in turn will make this method of question selection more user friendly. This coupled with the strength classification characteristics of the DT and the use of Bayesian updating will result in a robust method to use in question selection that will ensure high probability of acceptance.

Chapter 3
Acceptance Model

3.0 Introduction

Profit comes from good accounts but this can only materialise when a good customer accepts an offer from the organisation. So it is also necessary to get the customers to accept the offer.

In order to "persuade" a customer to accept an offer, the lender must have information on the preferences of the customer as a guide on the type of offer that may be of interest. A lender can "learn" about the customers' preferences by looking at which type of product different types of customers accepted and hence has to decide what offer to make. In this model of the acceptance problem, the lenders decision problem on which offer to make is modelled as a Markov Decision Process under uncertainty.

The aim of this chapter is to develop an acceptance model using adaptive dynamic programming where Bayesian updating methods are employed to better estimate a take-up probability distribution. The significance of Bayesian updating in this model is that it allows previous responses to be included in the decision process. This means one uses learning of the previous responses to aid in selecting offers best to be offered to prospective customers that ensure take-up.

The standard connection between optimisation and mining the data in data warehouses is how to use the data to optimise the operations of the organisation. In this chapter, the focus is on an earlier problem in the process, namely what is the optimal way of collecting the data needed. This is done in the context where one acquires the data so as to estimate the propensity of customers to purchase the different variants of a product. The objective is to do this in such a way that optimises the overall profitability. Thus not only does one need to be optimal in the use of data, but also optimal in the way one acquires it.

Hence one has to look at how organisations learn the behaviour of customers who are making decisions on whether to accept the variant of a generic product that they are offered. This information is used to develop a model that utilises Bayesian updating methods to estimate the take-up probability distributions. The model can then be used to aid in selecting the offers to make to prospective customers to ensure take-ups which maximise the organisation's profits.

3.1 The Problem

Retailers have a number of variants of a product which they can offer their customers. The variants can be ordered so their attractiveness to the customers and so the likelihood of the customer accepting that variant is monotonically decreasing, while the variant's profitability to the lender is monotonically increasing. If this is not the case, and a product has both lower acceptance rate and lower profitability than another, then as far as the order is concerned, it is dominated by the other both in profitability and market share. Then the former variant should never be offered. We were motivated by the example of a credit card which can charge different interest rates, as part of its risk-based pricing. The decision then is which offer to make to the next applicant, given knowledge of the previous offers and whether they were accepted, and the objective is to maximise the profit to the lender.

In the models considered here, we assume a homogenous population of customers, so that the chance any individual will select variant i is p_i and the retailer makes a profit of P_i from each customer using product i where $p_1 \geq p_2 \geq \ldots \geq p_n$ and $P_1 \leq P_2 \leq \ldots \leq P_n$ where $P_i \geq 0$. Thus it would seem obvious that the retailer chooses to offer the variant which maximises his expected profit i.e. $\max\{p_i P_i\}$. However, in reality, the retailer does not know the probabilities p_i, save that $p_1 \geq p_2 \geq p_3 \ldots \geq p_n$. Nor does he know exactly the number of potential customers to whom an offer can be made. To keep the model relatively simple we assume the number of potential customers has a geometric distribution with parameter β, so the chance the current customer is going to be the last is $1-\beta$. The objective is to maximise the profit over this random horizon of customers. Thus one is choosing which offer to make so as both to learn what the probabilities are of p_i and at the same time ensure the profit is as large as possible.

We concentrate first on the model with only two variants. In the later sections, it will be shown how this can be extended to the *N* variant one. This problem has strong similarities with bandit problems first introduced by Gittins (1989) and now widely used Gittins (1979). However, it is not a bandit problem because there is interaction between the arms.

The following assumptions are made in this model. We assume that if a customer rejected variant *i*, that means that he would have also rejected all worse variants $j, j > i$. Similarly if he accepted variant *i*, he would have accepted all better variants $j, j < i$. We ensure this by defining a set of conditional probabilities:

$q_1 = p_1$

q_2= Prob.(customer would accept variant 2| customer would accept variant 1)

Since $q_1 = p_1$,

hence $p_2 = q_2 p_1$.

This condition ensures that $p_1 \geq p_2$.

Since the q_i are all Bernoulli random variables, in a Bayesian setting, one could describe the retailer's knowledge of the information by a Beta distribution- the prior family. For q_i, we describe its prior by $B(r_i, n_i)$ whose density function is $q_i^{r_i-1}(1-q_i)^{n_i-r_i-1}$, and expectation ($r_i / n_i$) were r_i=number of customers that have accepted the offer i and n_i=number of customers who were extended offer i. At any point, the retailer's belief about the acceptance probabilities $p_1 \geq p_2$ is given by the parameters (r_1, n_1, r_2, n_2). Let $V(r_1, n_1, r_2, n_2)$ be the expected maximum total future profit to the retailer given that his current belief is (r_1, n_1, r_2, n_2).

r_1, n_1 = the parameters of the Beta distribution describing one's belief of p_1. If offer 1 is accepted, the parameters get updated to $r_1 + 1$, $n_1 + 1$, while when it is rejected they get updated to r_1, $n_1 + 1$. Thus one could reinterpret these as

r_1 = the number of people already accepted Offer 1; and

n_1 = the number of people already offered Offer 1.

r_2, n_2 = the parameters of the Beta distribution describing one's belief of p_2. If an offer 2 is accepted, the parameters get updated to $r_2 + 1$, $n_2 + 1$, while it is rejected but the customer would have accepted Offer 1 they get updated to r_2, $n_2 + 1$ and the r_1, n_1 remain unchanged. Thus one could reinterpret these as

r_2 = the number of people already accepted Offer 2, given they would accept Offer 1, and

n_2 = the number of people who have been offered Offer 2 who would have accepted 1.

In both of the cases, observe that $n_i \geq r_i$ for $i = 1, 2, 3, \ldots$.

The selection of which product to offer is done by referring to information obtained based on past acceptance and rejections of the products. It is considered to be a "learning" model as it is basing its decision of which product to present to a prospective customer on past information.

With such belief distribution, the expected probability of product 1 being accepted if it is offered is $\frac{r_1}{n_1}$ and of product 2 is $\frac{r_1}{n_1}\frac{r_2}{n_2}$. Let $V(r_1, n_1, r_2, n_2)$= expected maximum future profit from the next customer. Given that one has to choose which of the 2 variants of the product to offer to the next customer, function $V(r_1, n_1, r_2, n_2)$ satisfies the optimality equation (Puterman, 1994).

$$V(r_1,n_1,r_2,n_2) = \max\begin{cases} \frac{r_1}{n_1}P_1 + \beta\{\frac{r_1}{n_1}V(r_1+1,n_1+1,r_2,n_2) + \left(1-\frac{r_1}{n_1}\right)V(r_1,n_1+1,r_2,n_2)\}; \\ \frac{r_1}{n_1}\frac{r_2}{n_2}P_2 + \beta\{\frac{r_1}{n_1}\frac{r_2}{n_2}V(r_1+1,n_1+1,r_2+1,n_2+1) + \\ \frac{r_1}{n_1}\left(1-\frac{r_2}{n_2}\right)V(r_1+1,n_1+1,r_2,n_2+1) + \left(1-\frac{r_1}{n_1}\right)V(r_1,n_1+1,r_2,n_2)\}. \end{cases} \quad (1)$$

This first term in each offer is the probability that the next customer will accept variant 1 or variant 2, respectively, multiplied by the profit to the lender. The remaining terms depends on the chance β that there will be another customer. The first corresponds to the current offer being accepted and so q_1 and q_2 all being "successful" with the customer. The remaining terms correspond to the offer being refused and it looks at the different ways it can happen. The term $V(r_1+1, n_1+1, r_2, n_2+1)$ corresponds to the refusal of the variant 2 where one believes the first variant that would have been refused is variant 2 while $V(r_1, n_1+1, r_2, n_2)$ means one believes variant 1 would also have been refused.

3.2 Optimal Solution for the Two-Variant Case

Consider a variation of the problem in **(1)** where the retailer has a cost of $\beta\left(1-\frac{r_1}{n_1}\right)V(r_1, n_1+1, r_2, n_2)$ if an offer is made to a customer where the state is (r_1, n_1, r_2, n_2) irrespectively of which offer is made. Since the cost is independent of the offer made, it cannot affect the optimal action. Let $\tilde{V}(r_1, n_1, r_2, n_2)$ be the optimal expected profit for the modified problem. Then, we know the optimal policy when solving for $\tilde{V}(r_1, n_1, r_2, n_2)$ is the same as for $V(r_1, n_1, r_2, n_2)$, with

$$\tilde{V}(r_1, n_1, r_2, n_2)$$

$$= \frac{r_1}{n_1}\beta \max \begin{cases} B_1 + \tilde{V}(r_1+1, n_1+1, r_2, n_2), \\ \frac{r_2}{n_2}B_2 + \frac{r_2}{n_2}\tilde{V}(r_1+1, n_1+1, r_2+1, n_2+1) + \left(1-\frac{r_2}{n_2}\right)\tilde{V}(r_1+1, n_1+1, r_2, n_2+1), \end{cases}$$

(2)

where $B_i = \frac{P_i}{\beta}$ and $P_i \geq 0$.

Note that in **(2)** r_1, n_1 only appear in the right hand side either outside the maximisation or in the form $\tilde{V}(r_1+1, n_1+1, r_2, n_2)$. In the latter terms the values of r_1, n_1 both go up by one each time an offer is made irrespective of the outcome. This means we can simplify the optimality equation. Assume the initial belief of q_1 is given by a Beta distribution with the parameters R, N. With this choice of parameterization, after *m* offers are made which has resulted in *r* successes, R, N will be updated to $(R+r, N+m)$ which we represent in general as r_1, n_1. We then define a new set of discount factors by $\beta_s = \frac{\beta(R+s)}{N+s}$ so that if we define

$$W_s(r_2, n_2) = \tilde{V}(R_1+s, N_1+s, r_2, n_2),$$

then $W_s(r_2, n_2)$ satisfies

$$W_s(r_2, n_2) = \beta_s \max\{B_1 + W_{s+1}(r_2, n_2); \frac{r_2}{n_2}B_2 + \frac{r_2}{n_2}W_{s+1}(r_2+1, n_2+1) + \left(1-\frac{r_2}{n_2}\right)W_{s+1}(r_2, n_2+1)\}$$

for all $n_2 = 1, 2, ..., r_2, \quad 0 \leq r_2 \leq n_2, \quad s = 0, 1, 2, ...$

The advantage of working with this form of the optimality equation is that the impact of the belief about q_i has been reduced to the subscript *s* of the discount factor. Now we can prove some general results.

Lemma 3.1

$W_s(r_2,n_2)$ is (i) non decreasing in r_2.

(ii) non increasing in n_2.

(iii) non decreasing in s.

Proof

These results are proved by using induction on the iterates of value iteration. Consider the value iteration scheme for all $m = 0,1,2,...$

$$W_s^{m+1}(r_2,n_2) = \beta_s \max\begin{cases} B_1 + W_{s+1}^m(r_2,n_2), \\ \frac{r_2}{n_2}B_2 + \frac{r_2}{n_2}W_{s+1}^m(r_2+1,n_2+1) + \left(1-\frac{r_2}{n_2}\right)W_{s+1}^m(r_2,n_2+1), \end{cases}$$

with $W_s^0(r_2,n_2) = 0 \;\; \forall s, r_2, n_2$.

Trivially $W_s^1(r_2,n_2) \geq W_s^0(r_2,n_2) \;\; \forall s, r_2, n_2$

and then we will use induction to prove

$W_s^{m+1}(r_2,n_2) \geq W_s^m(r_2,n_2) \;\; \forall s, r_2, n_2$.

Since $\max\{a,b\} - \max\{c,d\} \geq \min\{a-c, b-d\}$,

$W_s^{m+1}(r_2,n_2) - W_s^m(r_2,n_2) \geq \beta_s \min_s\{W_s^m(r_2,n_2) - W_s^{m-1}(r_2,n_2)\} \geq 0$.

Thus $W_s^{m+1}(r_2,n_2)$ is a monotone increasing sequence bounded above by $\left\{\frac{B_2}{\beta}\right\}$,

and so the iterates of value iteration converge to $W_s(r_2,n_2)$(see Puterman (1994)). To prove (i), (ii), and (iii), we use induction to prove the same results hold for the iterates of value iteration.

(i) $W_s^m(r_2+1,n_2) - W_s^m(r_2,n_2) \geq 0$.

Assume true for W_s^m and note

$$W_s^{m+1}(r_2+1,n_2)-W_s^{m+1}(r_2,n_2)\ge \beta_s \min\{W_{s+1}^m(r_2+1,n_2)-W_{s+1}^m(r_2,n_2);$$
$$\frac{1}{n_2}(B_2+W_{s+1}^m(r_2+2,n_2+1)-W_{s+1}^m(r_2+1,n_2+1)+\frac{r_2}{n_2}(W_{s+1}^m(r_2+2,n_2+1)-W_{s+1}^m(r_2+1,n_2+1))+$$
$$\left(1-\frac{r_2}{n_2}\right)(W_{s+1}^m(r_2+1,n_2+1)-W_{s+1}^m(r_2,n_2+1))\}$$

which is positive. If the hypothesis holds for m, then the hypothesis holds for $m+1$ and so holds in the limit.

Similarly, for (ii), one assumes

$W_s^m(r_2,n_2)-W_s^m(r_2,n_2+1)\ge 0$ and then one uses the following that

$$W_s^{m+1}(r_2,n_2)-W_s^{m+1}(r_2,n_2+1)\ge \beta_s \min\{W_s^m(r_2,n_2)-W_s^m(r_2,n_2+1);$$
$$\frac{r_2}{n_2(n_2+1)}(B_2+W_s^m(r_2+1,n_2+1)-W_s^m(r_2,n_2+2))+\frac{r_2}{(n_2+1)}(W_s^m(r_2+1,n_2+1)-$$
$$W_s^m(r_2+1,n_2+2))+\left(1-\frac{r_2}{n_2}\right)(W_s^m(r_2,n_2+1)-W_s^m(r_2,n_2+2))\}\ge 0.$$

(iii) The monotonicity in s follows in a similar way using the induction hypothesis $W_{s+1}^m(r_2,n_2)-W_s^m(r_2,n_2)\ge 0$.

When we are completely confident of our knowledge of q_i, the form of the optimal policy is obvious. This complete knowledge corresponds to the case when $s=\infty$ and the optimal policy is of a limit form where one chooses variant 1 if the belief is that q_2 is below a certain level and choose variant 2 if it is above this level. The results are formalized in the following theorem:

Theorem 3.1

In the limit when $s=\infty$, the optimal policy is of the form that there exists a function $r^*(n_2)$ so that:

(a) in (r_2,n_2) with $r_2\le r^*(n_2)$, one chooses variant 1.

(b) in (r_2,n_2) with $r_2> r^*(n_2)$, one chooses variant 2.

Moreover, $r^*(n_2)$ is non decreasing in n_2.

Proof

Consider the optimality equation

$W_\infty(r_2,n_2)=\beta\max\{W_\infty^1(r_2,n_2);W_\infty^2(r_2,n_2)\}$ where

$W_\infty^1(r_2,n_2)=B_1+W_\infty(r_2,n_2)$ and

$$W_\infty^2(r_2,n_2)=\frac{r_2}{n_2}B_2+\frac{r_2}{n_2}W_\infty(r_2+1,n_2+1)+\left(1-\frac{r_2}{n_2}\right)W_\infty(r_2,n_2+1). \qquad \textbf{(3)}$$

If for any (r_2,n_2), one chooses option 1 then

$W_\infty(r_2,n_2)=\beta(B_1+W_\infty(r_2,n_2))$ or

$W_\infty(r_2,n_2)=\dfrac{\beta B_1}{1-\beta}=W_\infty^1(r_2,n_2)$.

This value is independent of r_2 and n_2. For a given n_2, the fact $W_\infty(r_2,n_2)$ is non decreasing in r_2 (from Lemma 1) means that if r_2 chooses variant 1, since the first term on the right hand side is greater than the second, then for any $r\le r_2$, this inequality will continue to hold and variant 1 must be chosen. So,

$$r^*(n_2)=\max\left\{r_2:W_\infty(r_2,n_2)=\frac{\beta B_1}{1-\beta}=W_\infty^1(r_2,n_2)\right\},$$

$$\max\{r_2:W_\infty^1(r_2,n_2)\ge W_\infty^2(r_2,n_2)\}.$$

For all $r_2>r^*(n_2)$, $W_\infty^2(r_2,n_2)$ in the right hand side of **(3)** remains greater than $W_\infty^1(r_2,n_2)$ which ensures the optimal policy is to choose Variant 2. The fact that $W_\infty^2(r_2,n_2)$ is non increasing in n_2 while $W_\infty^1(r_2,n_2)$ is independent of both r_2 and n_2, implies that $r^*(n_2)$ is non decreasing in n_2.

Having proved the form of the optimal policy for the limiting case, when we have dealt with a large number of customers, we turn to the case where s customers have already been made an offer. Another transformation will make the results of this case more obvious. Define

$D_s(r,n)=W_s(r,n)-W_s(0,n)\quad\forall r,s,n$.

Note that $W_s(0,n)$ is the case where one believes $q_2 = 0$ and Bayes theorem says one will continue to believe this. Hence in this case, one always offers variant 1 and so in fact $W_s(0,n) = \sum_{r=0}^{\infty}\left(\prod_{k=0}^{r}\beta_{s+k}\right)B_1$, which is independent of n.

Note that $D_s(r,n)$ satisfies the optimality equation

$$D_s(r,n) = \beta_s \max\left(D_{s+1}(r,n); \left(\frac{r}{n}B_2 - B_1\right) + \left(\frac{r}{n}\right)D_{s+1}(r+1,n+1) + \left(\frac{n-r}{n}\right)D_{s+1}(r,n+1)\right). \quad \textbf{(4)}$$

One can then modify Lemma 3.1 trivially to obtain

Lemma 3.2

$D_s(r_2,n_2)$ is non decreasing in r_2 and s, and non increasing in n_2.

Proof

The monotonicity in r_2 and n_2 follow directly from Lemma 3.1. To prove the monotonicity in s, we need to use induction on the iterates of value iteration as in Lemma 3.1. Recall from the proof of Theorem 3.1 that if $D_s(r,n) = 0$, then the optimal policy is always to offer variant 1 thereafter, since $W_s(r,n)$ and $W_s(0,n)$ have the same value and so must use the same policy.

If $D_s(r,n) > 0$, then the optimal policy must involve offering variant 2 to some future customers. This leads to the following

Theorem 3.2

There exists a function $r^*(n)$ so that for any state (s, r_2, n_2) having offered to s customers and having a belief about q_2 given by Beta(r_2, n_2), then if

(i) $r_2 \le r^*(n_2)$, offer variant 1 to all future customers.

(ii) $r_2 > r^*(n_2)$, variant 2 will be offered to at least one future customer.

Proof

Consider the case when $s=\infty$, then Theorem 3.1 defines a $r^*(n)$ so that for $r_2 \le r^*(n_2)$ in (r_2, n_2), offer Variant 1 to the current and hence all future customers. This means $D_\infty(r_2, n_2)=0$. For any other s, Lemma 2 implies $D_s(r, n) \le D_\infty(r, n)$. Hence for $r_2 < r^*(n_2)$, $D_s(r_2, n_2) \le D_\infty(r_2, n_2)=0$ and so the optimal policy is the offer variant 1 to all future customers.

If $r_2 > r^*(n_2)$, then in (r_2, n_2), one can offer variant 2 and so $D_\infty(r_2, n_2)>0$. If for a fixed n_2, $r_2 > r^*(n_2)$, then $D_\infty(r_2, n_2)>\varepsilon$. Then by convergence of $W_s(r_2, n_2)$, there exists a t^* so that for $t \ge t^*$, $D_t(r_2, n_2)>\frac{\varepsilon}{2}$. Then

$$D_s(r_2, n_2) \ge \beta_s \beta_{s+1} \ldots \beta_{t^*} \frac{\varepsilon}{2}>0.$$

So for $r_2 > r^*(n_2)$, one will offer variant 2 to some future customers whatever the number s of customers to whom offers have already been made.

Corollary

If $r_2 > \frac{B_1}{B_2} n_2$, then one offers variant 2 to the next customer whatever s is.

Proof

Assume the optimal policy is π_1 where, $\pi_1=(1)^k(2)\tilde{\pi}$, that is Offer Variant 1 to the next k customers and then Variant 2 to the $k+1^{ith}$ and any policy thereafter. We show that the policy cannot be optimal if $\frac{r_2}{n_2} B_1 > B_2$, by comparing it with policy π_2 where, $\pi_2=(2)(1)^k\tilde{\pi}$,

$$D_s^{\pi_2}(r_2, n_2)-D_s^{\pi_1}(r_2, n_2)=\beta_s\left(1-\prod_{r=1}^{k} \beta_{s+r}\right)\left(\frac{r_2}{n_2} B_1-B_2\right)>0.$$

If one were to choose variant 1 in state (s, r_2, n_2), the policy must be of the form $(1)^k(2)\tilde{\pi}$, for given k, which is less than $(2)(1)^k\tilde{\pi}$. Hence one must choose variant 2.

Thus, we have shown that the optimal policy is such that there exists a $r^*(n)$ so that in state (s, r_2, n_2), if

(i) $r_2 \le r^*(n_2)$, one chooses variant 1 for all future offers.

(ii) $r_2 \ge \frac{B_1}{B_2} n_2$, one chooses variant 2.

In some cases, there is no gap between theses two inequalities but in other cases, there may be. In that situation, we conjecture there is a function $r^*(s, n)$ so that for $r_2 \le r^*(s, n_2)$, one chooses variant 1 and otherwise one chooses variant 2. We will assume this form of the policy in presenting the results of our numerical calculations in the next section.

3.3 Experimental Results

There is no publicly available data on the offer strategies used by the personal financial sector. Discussion with practitioners suggests that in general, they offer the same product to all people and the choice of offer is usually subjective. If they were to look at the data for such offers, the profit maximizing strategy would be to choose variant 1 if $\frac{R_1}{N_1} P_1 > \frac{R_1}{N_1} \frac{R_2}{N_2} P_2$ and variant 2 otherwise, that is choose variant 2 if $R_2 < N_2 \frac{P_1}{P_2}$.

In calculating the following results, we used the work on bounds on the optimal value function in Markov decision processes by MacQueen (1966) and Porteus (2002). They obtained the upper and lower bounds on the optimal solution. These bounds were then utilised to decide and to recognise when to stop the calculation because the difference between the upper and lower bounds was within an acceptable tolerance.

The following tables contain examples of results generated by the model for the case when we start with $\beta = 0.5$ and apply to the model for $\widetilde{V}$. Recall these are not the full profits because we have subtracted the 'fee' from them. The policies are the ones that are optimal for the original problem. All the $r^*(n_2)$ in the following examples have been calculated offline. This then allows us to make offer decisions in real time by looking at the current values of the parameters.

If our initial beliefs were $R_1 = 1, N_1 = 2, R_2 = 1, N_2 = 2$ (which is the uniform distribution for q_1 and q_2) the resulting optimal strategy decision, would be a profit of 12.5 per applicant.

We represent some of the belief points at which the offer decision changes in **Tables 3.1** and **3.2**. Note the bold results are when the switch of offers occurs. We choose $r_1 = 6, n_1 = 16$ to represent a case where one's belief of the acceptance of variant 1 is fairly tightly spread around $p_1 = \frac{3}{8}$.

($P_1 = 10.000, P_2 = 25.000, \beta = 0.5$)

r_1	n_1	r_2	n_2	**Profit (£)**	**Decision**
6	16	0-1	3	4.7481	1
6	**16**	**2**	**3**	**7.9144**	**2**
6	16	3	3	11.8702	2

Table 3.1: Part of results generated by model

($P_1 = 10.000$, $P_2 = 25.000$, $\beta = 0.5$)

r_1	n_1	r_2	n_2	**Profit (£)**	**Decision**
6	16	0-9	24	4.7481	1
6	**16**	**10**	**24**	**4.9490**	**2**
6	16	11	24	5.4406	2
6	16	12	24	5.9351	2
6	16	13	24	6.4297	2
6	16	14	24	6.9243	2
6	16	15	24	7.4189	2
6	16	16	24	7.9135	2
6	16	17	24	8.4081	2
6	16	18	24	8.9027	2
6	16	19	24	9.3973	2
6	16	20	24	9.8919	2
6	16	21	24	10.3865	2
6	16	22	24	10.8811	2
6	16	23	24	11.3757	2
6	16	24	24	11.8702	2

Table 3.2: **Part of results generated by model**

One might wonder if it is the mean of one's belief that matters or does the variance of q_1 also affect the decision. This translates into whether it is only the ratio of r_1 to n_1 and not their values that matter.

In most cases including the one in **Table 3.3** and **Table 3.4** this is true. In these tables, we define for what values of n_2, $r^*(n_2) = r_2$.

$r_1 = 4, n_1 = 5$

r_2	0	1	2	3	4	5	6	7	8	9	10	11
n_2	1-3	4-6	7-8	9-11	12-13	14-16	17-18	19-21	22-23	24-26	27-28	29-30

Table 3.3: **Changing of offers when** $r_1 = 4, n_1 = 5, P_1 = 20.000, P_2 = 50.000, \beta = 0.999$

$r_1 = 16, n_1 = 20$

r_2	0	1	2	3	4	5	6	7	8	9	10	11
n_2	1-3	4-6	7-8	9-11	12-13	14-16	17-18	19-21	22-23	24-26	27-28	29-30

Table 3.4: **Changing of offers when** $r_1 = 16, n_1 = 20, P_1 = 20.000, P_2 = 50.000, \beta = 0.999$

However, there are counter examples to the hypothesis, namely in **Tables 3.5** and **3.6**. Note the bold numbers indicate the where for the values of n_2, $r^*(n_2) \neq r_2$.

$r_1 = 2, n_1 = 3,$

r_2	0	1	2	3	4	**5**	**6**	**7**	**8**	**9**	**10**	11
n_2	1-3	4-5	6-8	9-10	11-13	**14-16**	**17-18**	**19-21**	**22-23**	**24-26**	**27-28**	29-30

Table 3.5: Changing of offers when $r_1 = 2, n_1 = 3, P_1 = 20000, P_2 = 50000, \beta = 0.999$

$r_1 = 16, n_1 = 24$

r_2	0	1	2	3	4	**5**	**6**	**7**	**8**	**9**	**10**	11
n_2	1-3	4-5	6-8	9-10	11-13	**14-15**	**16-18**	**19-20**	**21-23**	**24-25**	**26-28**	29-30

Table 3.6: Changing of offers when $r_1 = 16, n_1 = 24, P_1 = 20000, P_2 = 50000, \beta = 0.999$

Thus at $r_1 = 16, n_1 = 24$, the switch of offers occurs later, at $r_2 = 7, n_2 = 16$, than in the case $r_1 = 2, n_1 = 3$, when the switch occurs at $r_2 = 6, n_2 = 16$.

Tables 3.7, **3.8** and **3.9** below show that as r_1 increases but the rate $\frac{r_1}{n_1}$ is fixed, the crucial value where one changes offers, $r^*(n_2)$, is monotonically non increasing. This is not surprising as Variant 1 is becoming a better bet and so one is willing to start using it, when previously one would want to experiment with Variant 2.

We also give results for the effect of more information (increment of r_1) in the tables. In each table, we report $r^*(n_2)$ for n_2 varying from 1 to 30.

r_1, n_1	n_2	1	2	3	4	5	6	7	8	9	10
1, 24	$r_2^*(n_2)$	0	0	1	1	1	2	2	3	3	3
4, 24		0	0	1	1	1	2	2	3	3	3
8, 24		0	0	1	1	1	2	2	3	3	3
16, 24		0	0	0	1	1	2	2	2	3	3
22, 24		0	0	0	1	1	1	2	2	2	3
23, 24		0	0	0	0	1	1	1	2	2	2
24, 24		0	0	0	0	0	1	1	2	2	2

Table 3.7: Effect of more information on the switch of offers

r_1, n_1	n_2	11	12	13	14	15	16	17	18	19	20
1,24	$r_2^*(n_2)$	4	4	5	5	5	6	6	7	7	7
4, 24		4	4	5	5	5	6	6	7	7	7
8, 24		4	4	5	5	5	6	6	7	7	7
16, 24		4	4	4	5	5	6	6	6	7	7
22, 24		3	4	4	4	5	5	5	6	6	7
23, 24		3	3	4	4	4	5	5	6	6	6
24, 24		3	3	3	4	4	4	5	5	6	6

Table 3.8: Effect of more information on the switch of offers

r_1, n_1	n_2	21	22	23	24	25	26	27	28	29	30
1, 24	$r_2^*(n_2)$	8	8	9	9	9	10	10	11	11	11
4, 24		8	8	9	9	9	10	10	11	11	11
8, 24		8	8	9	9	9	10	10	11	11	11
16, 24		8	8	8	9	9	10	10	10	11	11
22, 24		7	7	8	8	9	9	9	10	10	11
23, 24		7	7	7	8	8	9	9	9	10	10
24, 24		6	7	7	8	8	8	9	9	10	10

Table 3.9: Effect of more information on the switch of offers

3.4 Many Variant Problem

Where there are N variants of the product, we can extend the approach used for the 2 variant problem as follows:

As in the 2 variant case, we assume that if a customer rejected variant i, all worse variants $j, j > i$ will also be rejected. Similarly if variant i is accepted, all better variants $j, j < i$ will be accepted. This is ensured by defining a set of conditional probabilities:

$q_1 = p_1$,

q_{i+1} = Prob (customer would accept variant i+1 | customer would accept variant i).

Since $q_1 = p_1$, hence $p_2 = q_2 p_1$.

So $p_{i+1} = q_{i+1} p_i = \prod_{j=1}^{i+1} q_j$.

This condition ensures that $p_1 \geq p_2 \geq p_3 ... \geq p_n$.

Since the q_j are all Bernoulli random variables, in a Bayesian setting, one would describe the retailer's knowledge of the information (the prior family) by a Beta distribution. For q_i, we describe its prior by $B(r_i, n_i)$ with a density function of $q_i^{r_i - 1}(1 - q_i)^{n_i - r_i - 1}$, and expectation ($r_i / n_i$). At any point, the retailer's belief about the acceptance probabilities $p_1 \geq p_2 \geq p_3 ... \geq p_n$ is given by the parameters $(r_1, n_1, r_2, n_2, ..., r_N, n_N)$. So, let $V(r_1, n_1, r_2, n_2, ..., r_N, n_N)$ be the expected maximum total future profit to the retailer given that his current belief is $(r_1, n_1, r_2, n_2, ..., r_N, n_N)$.

With such belief distribution, the expected probability of product k being accepted if it is offered is $\prod_{j=1}^{k} \frac{r_j}{n_j}$. Given that one has to choose which of the N variants of the product to offer to the next customer, function $V(r_1, n_1, r_2, n_2, ..., r_N, n_N)$ satisfies the optimality equation.

$$V(r_1,n_1,r_2,n_2,\dots,r_N,n_N)=\max_{1\le k\le N}\{\prod_{j=1}^{k}\left(\frac{r_j}{n_j}\right)P_k+\beta(\prod_{j=1}^{k}\left(\frac{r_j}{n_j}\right)V(r_1+1,n_1+1,\dots,r_j+1,n_j+1)+$$

$$\sum_{l=1}^{k}\left(\prod_{j=1}^{r-1}\left(\frac{r_j}{n_j}\right)\left(1-\frac{r_l}{n_l}\right)\right)V(r_1+1,n_1+1,\dots,r_l,n_l+1))\}. \qquad \textbf{(5)}$$

This first term is the probability that the next customer will accept variant k multiplied by the profit to the lender. The remaining terms depends on the chance β that there will be another customer. The first corresponds to the current offer being accepted and so $q_1,q_2,\dots,q_l$ all being "successful" with the customer. The remaining terms correspond to the offer being refused and it looks at the different ways it can happen. The term $V(r_1+1,n_1+1,\dots,r_l,n_l+1)$ corresponds to the refusal of the *k*th variant where one believes the first variant that would have been refused is the *l*th one.

Recall that **(5)** is the optimality equation for N variants of the product which is the extension of **(1)** in the 2-variant case. Subtract a cost of $\left(1-\frac{r_1}{n_1}\right)V(r_1,n_1+1,r_2,n_2,\dots,r_N,n_N)$ from all the actions in state $(r_1,n_1,\dots,r_N,n_N)$ of **(4)**. This cannot affect the decisions made but allows us to simplify **(4)** to the equation

$$\tilde{V}(r,n_1,\dots,r_N,n_N)=\beta\max_{R\le 1\le k\le N}\{\prod_{j=1}^{k}\frac{r_j}{n_j}B_k+\prod_{j=1}^{k}\frac{r_j}{n_j}+V(r_1+1,n_1+1,\dots,r_k+1,n_k+1,r_{k+1},n_{k+1},\dots)+$$

$$\sum_{r=2}^{k}\left(\prod_{j=1}^{r-1}\frac{r_j}{n_j}\left(1-\frac{r_m}{n_m}\right)V(r_1+1,n_1+1,\dots,r_m,n_m+1,\dots,r_N,n_N)\right)\}. \qquad \textbf{(6)}$$

If its initial beliefs about the probability q_1 are given by R_l,N_i, then after s offers are made, the state parameters describing q_1 in this modified problem will be (R_1+s,N_1+s). If we define

$W_s(r_2,n_2,...,r_N,n_N) = \widetilde{V}(R_1+s, N_1+s, r_2, n_2...)$ and $\beta_s = \dfrac{R_1+s}{N_1+s}\beta$, then we have an optimal equation essentially independent of r_1, n_1 namely

$$W_s(r_2,n_2,...,r_N,n_N) = \beta_s \max_{k:1\le k\le N} \{\prod_{j=2}^{k}\frac{r_j}{n_j}(B_k + W_s(r_2+1,n_2+1,...,r_k+1,n_k+1,r_{k+1},n_{k+1},...)) +$$

$$\sum_{m=2}^{N}\left(\prod_{j=1}^{r-1}\frac{r_j}{n_j}\left(1-\frac{r_m}{n_m}\right)V(r_1+1,n_1+1,..,r_m,n_m+1,r_m,n_{m+1},...)\right)\}.$$

Lemma 3.3

$W_s(r_2,n_2,..,r_N,n_N)$ is (i) non decreasing in r_i.

(ii)non increasing in n_i.

(iii)non decreasing in s.

This leads to the case with no extra information on q_1, namely

Theorem 3.3

In the case when $s=\infty$, the optimal policy is of the form that there exists at function $r^*(r_2,n_2,...,r_N,n_N)$ so that in $(r_2,n_2,...,r_N,n_N)$:

(a) with $r_2 \le r^*(r_2,n_2,...,r_N,n_N)$, one chooses variant 1

(b) with $r_2 > r^*(r_2,n_2,...,r_N,n_N)$, one chooses some other variant other than variant 1

One can get results for the general s case in the same way as was done for the 2-variant problem by defining

$D_s(r_2,n_2,...,r_N,n_N) = W_s(r_2,n_2,...,r_N,n_N) - W_s(0,n_2,...,0,n_N)$

and proving again that $D_s(r_2,n_2,...,r_N,n_N)$ is non decreasing in r_2 and non increasing in n_2.

It then follows that if for any $(r_2,n_2,...,r_N,n_N)$, $D_s(r_2,n_2,...,r_N,n_N)=0$, then one should choose Variant 1 and in fact continue to choose it. Similarly, one can show using the argument in Theorem 3.2 that if $D_s(r_2, n_2, ..., r_N, n_N) > 0$, one will choose to offer some other variant than 1 to some customer in the future. Hence we have the result

Theorem 3.4

For every s, there exists a function of $r_s^*(n_2, ..., r_N, n_N)$ so that the optimal policy in state (r_2, n_2) is to choose variant 1 for all future customers if $r_2 \le r^*(n_2, ..., r_N, n_N)$ and to offer some other variant to some of the future customers if $r_2 > r^*(n_2, ..., r_N, n_N)$.

Once one has defined the region where one would choose variant 1, one can then concentrate on the regions where variants 2 to N should be chosen. The whole process can then be repeated on the region ignoring variant 1. This leads to the description of the optimal policy as follows:

Theorem 3.5

At any state $(r_1, n_1, ..., r_N, n_N)$, there exists functions

$r_i^*(r_{i+1}, n_{i+1}, ..., r_N, n_N)$, $i = 2,3,...,N$ so that

(i) one offers variant 1 to all future customers if $r_2 \le r_2^*(n_2, ..., r_N, n_N)$; if variant 1 is not chosen then.

(ii) one offers variant 2 to all future customers if $r_3 \le r_3^*(n_3, r_4, n_4, ..., r_N, n_N)$; if variants, 1,2 , ..., $s-1$ are not chosen then.

(iii) one chooses variant s to all future customers if $r_{s+1} < r_{s+1}^*(n_{s+1}, r_{s+2}, n_{s+2}, ..., r_N, n_N)$; and continue the nesting implicit above.

We can use the optimality equation to compute solutions for the N variant problem. **Tables 3.10** and **3.11** gives the results for a 3-variant problem (for some sample of r_2, n_2). The profits here are the actual profits where the "fee" has not been taken. Note that the bold numbers correspond when the switch of offers occur.

$(P_1 = 20.000, P_2 = 50.000, P_3 = 80.000)$

r_1	n_1	r_2	n_2	r_3	n_3	**Profit (£)**	**Decision**
6	9	4	9	0	10	29.6296	2
6	9	4	9	1	10	29.6296	2
6	9	4	9	2	10	29.6296	2
6	9	4	9	3	10	29.6296	2
6	9	4	9	4	10	29.6296	2
6	9	4	9	5	10	29.6296	2
6	9	4	9	6	10	29.6296	2
6	**9**	**4**	**9**	**7**	**10**	**33.1852**	**3**
6	9	4	9	8	10	37.9259	3
6	9	4	9	9	10	42.6667	3
6	9	4	9	10	10	47.4074	3

Table 3.10: Part of results generated by model for 3 variants

$(P_1 = 20.000, P_2 = 50.000, P_3 = 80.000)$

r_1	n_1	r_2	n_2	r_3	n_3	**Profit (£)**	**Decision**
6	9	3	9	0	10	26.6667	1
6	9	3	9	1	10	26.6667	1
6	9	3	9	2	10	26.6667	1
6	9	3	9	3	10	26.6667	1
6	9	3	9	4	10	26.6667	1
6	9	3	9	5	10	26.6667	1
6	9	3	9	6	10	26.6667	1
6	9	3	9	7	10	26.6667	1
6	**9**	**3**	**9**	**8**	**10**	**28.6933**	**3**
6	9	3	9	9	10	32.0000	3
6	9	3	9	10	10	35.5556	3

Table 3.11: Part of results generated by model for 3 variants

3.5 Conclusion

We have looked at the problem of which variant of a product to offer to a customer, so as to learn how the acceptance rate depends on the variant of the product offered. In learning, we also try to maximise the expected profit to the organisation. The problem was motivated by using the internet for credit card applications, where the credit card company can offer various interest rates or various credit limits but does not know the acceptance rate for each type of offer. The model assumes the population is homogeneous in its acceptance rate and likely profitability.

For a more realistic model, one could construct segments with different prior beliefs of acceptance, different profitability functions in each segment, and hence different offer policies for each segment. One can use the information gathered on the client during the application process to identify which segment a customer is likely to be in.

Hence it is clear that in future, companies will need to be clearer not only in the way they mine the data in their data warehouses, but also in the way they collect their data. In the offering process, one is trying to acquire data on how acceptance rates depend on the type of customer and the variant of the product offered, but in such a way that one is also making decisions that are as profitable as they can be.

Chapter 4

TAROT (Top Applicant characteristics Remainder Offer characteristics Tree)

4.0 Introduction

In Chapter 3, we developed an acceptance model which is able to predict which offer to make to the next customer. Before this decision can be made, we need to know what has been offered to which customer and whether they accepted the offers or not. Hence, another way of getting information on acceptance and non-acceptance of offers is to use past information analysed at one time. In particular, one can use this to find out not which offers to make, but which questions to ask of an applicant and then which offer to make. Once the data on past offers is collected, we can use techniques to perform the classification.

Classification and Regression Trees (CART) or decision trees are one of the techniques used in the classification of data. It was developed by Breiman (Breiman *et al.*, 1984) and remains a powerful classification tool, especially with the new advances like random forests in its use (Breiman, 2001). In classification one has a set of data, where each data point consists of the values (the attributes) of a number of characteristics variables, and a target variable, which is often binary. The aim is to identify which combinations of attributes lead to the different target outcomes.

A decision tree performs its classification in a top-down fashion; starting from the root node. From there, it works itself downwards based on the outcomes of the test at the internal nodes until a leaf node is achieved, where it then assigns a class label to that particular leaf node (Baesens, 2003). Hence, decision trees have a unique characteristic which makes them different from the rest of the classification techniques; it allows for the analysis of the interaction between the variables. It allows the user to "experiment" with different competing nodes and analyse the effect of such a selection on the leaf node. This proves useful when building a tree with certain specifications, in this case to achieve the highest acceptance rate of offers.

Decision trees are normally used to get the best classification tree, a structure that highlights the variables that best classify the data set at hand. When building a decision tree using SAS Enterprise Miner 4.3, one has a choice of splitting rules, for example chi-square (χ^2) or entropy. In order to better describe how the chi-square statistic is used in classification, we take the explanation from Thomas *et al.*, (2002). Let's first define

g_i = the numbers of good customers with attribute i and

b_i = the number of bad customers with attribute i

g and b are the total number of good customers and bad customer respectively in the population.

And

$$\hat{g}_i = \frac{(g_i + b_i)g}{g+b},$$

$$\hat{b}_i = \frac{(g_i + b_i)b}{g+b}$$

where $\hat{g}_i$ and $\hat{b}_i$ are the expected number of good customers and bad customers with attribute i given the ratio for attribute i is the same as for the whole population. The chi-square statistic is defined as

$$s^2 = \sum_i \left(\frac{(g_i - \hat{g}_i)^2}{\hat{g}_i} + \frac{(b_i - \hat{b}_i)^2}{\hat{b}_i} \right),$$

and is usually used to measure the likelihood of no differences in the good: bad ratio in the different classes. One compares it with the chi-square statistics with $k-1$ degrees of freedom, where k is the number of classes of the characteristic in question.

However the chi-square statistic can also be used as a measure of how different the odds are in the different classes, with a higher value reflecting greater differences in the odds. Thus, the larger the value of the chi-square statistic, the better the split hence the better classifier it is.

Entropy is the other splitting criterion considered when building a decision tree. It is an impurity index used to assess how impure each node v of a tree is, where a pure node is one consisting of only one class. The index is calculated as follows:
If $i(v) =$ entropy index for node v, then

$$i(v) = -p\,(\text{Good}|v)\log_2\left(p(\text{Good}|v)\right) - p\,(\text{Bad}|v)\log_2\left(p(\text{Bad}|v)\right).$$

with $p(\text{Good}|v) + p(\text{Bad}|v) = 1$.

The index measures the number of methods that a good: bad split can be obtained in a node. So essentially, the entropy index measures the order or even disorder (depending on objective) of the classes in the data. If $p(\text{Good}|v) = p(\text{Bad}|v) = 0.5$ (situation of maximum disorder, minimum order), then the entropy index will be at its maximum which is the value 1. If $p(\text{Good}|v) = 0$ or $p(\text{Bad}|v) = 0$ (situation of minimum disorder, maximum order), then the entropy index will be at its lowest value, 0. So, the splitting rule would use the node with the highest decrease in weighted entropy.

In this chapter, we use entropy as our splitting criterion. At each decision node, each variable will be considered and the splitting rule is used to select a split that produces the highest decrease in entropy. Given if a split is forced in at a decision node, the decision trees will continue to use the splitting criterion on the remainder of the tree.

At each decision node, note that the splitting rule (chi-square or in this case, entropy) used by the decision tree will split the data in a decision node into two subsets. It chooses a split of variables which results in the highest decrease in weighted entropy. As the target variable is the acceptance of an offer, the split that is chosen will be one that best separates those who have accepted an offer and those who have rejected the offer; with the accepts set having the highest acceptance rate amongst all the variables considered at that node.

In the process of building the decision tree, there will (might) be instances where a certain variable will be forced in, meaning the split has to be made on it. The tree will continue using the chosen splitting rule on that variable as well. This is an essential feature of the technique used in this chapter when trying to build a decision tree to help decide what next offer to extend to a customer, by asking only 2 to 3 questions.

The decision trees in this chapter are built with a target variable which is the action of the applicants accepting the offer made. The decision tree uses the interactions between the applicant variables to find the best splits available for the variable by choosing one that results in the best offer acceptance rate set for a specific offer characteristic-based leaf node. The objective of building this particular decision tree is to help find an offer which will have the highest acceptance rate. Then, we can use the applicant characteristics of that particular offer to formulate questions. We will use these questions to identify which group a customer belongs to and which offer to extend. The key is to achieve this by asking as few questions as possible.

Decision trees are used in many areas of decision making. Its success in health can be seen with CART analysis being used in deciding treatment that was both cost conscious and appropriate in the treatment of diabetic retinopathy (Harper *et al.*, 2003). Decision trees were again the chosen classification technique when Harper and Winslett (2006) successfully used decision trees to develop a model which helped doctors detect possible complications for pregnancies so as to decide on a possible course of treatment at birth.

This technique is so robust that it can be used in the development of an effective algorithm and heuristic for a generalized assignment problem (Haddadi and Ouzia, 2004) solving scheduling problems using decomposition (Al-Khayyal *et al.*, 2001) or even studying the behavioural and procedural consequence of structural variation (Pöyhönen *et al.*, 2001). In fact, examples like Brugha (2004), Ahn *et al.* (2000) and Ekárt and Németh (2005) all show how effective and popular decision trees are in decision making.

The work referenced in the paragraph before use decision trees without any or much modification. There was no structure imposed when building the decision tree. However, the decision tree that is going to be used in this chapter is not an ordinary classification tree. It is a deliberate attempt to structure a decision tree where the initial splits are on applicant characteristics which will produce leaf nodes consisting of offers to make to people that have the identified applicant characteristics.

This particular decision tree will be called TAROT (Top Application characteristics Remainder Offer characteristics Tree). The decision tree that is proposed here starts from classification based on only applicant characteristics, working from top to bottom, ending with an offer characteristic at each leaf node. Also, TAROT will allow splits to be forced into it. Hence it is not an optimal classification, but a classification based on a combination of applicant and offer characteristics.

Taking into account these requirements, there was a need for a statistical software package that has the ability to interactively train a decision tree with certain degree of ease. SAS 9.1.3. has an additional add-on called Enterprise Miner 4.3 which allows the user to experiment with "forced splits" and analyse its effects on the acceptance rates in the decision trees. This is achieved quite simply as the software displays the accepts and rejects in percentage form at each decision node on TAROT. Note that each of the competing nodes at each decision node would have already been tested and found to be of significant value. In this case, the entropy test is chosen and conducted on each applicant and offer variable. The other alternative would be software that is specifically developed to allow such "forced" splits, called SUPERCART 5.1 by Salford Solutions.

4.1 Data – Fantasy Student Account

There is no publicly available data on the offer strategies of the personal financial sector. So we utilise the data gathered from a Fantasy Student account which was started in 2001. The Fantasy Student account is a method of getting data on the acceptance of a financial product. It is an online website which enables students to apply for a fantasy current account. The website is accessible via this link (http://www.management.soton.ac.uk/staff/fairisaacs/). It consists of 3 pages which resemble the application forms for student accounts used by major banks in the UK. Please refer to the Appendix for a copy of the application form used. The questions created were based on questions used in application forms of 10 UK lenders, including 4 major UK retail banks. The questions on the first page were used to get data on demographic and financial information as well as interests. The student was then made an offer and the features of the offer would be outlined on the second page. There were all together 6 features of the account that could be varied. The student was then required to decide whether to accept or reject the offer (on the third page) and submit their decision.

The website was made available to everyone at the University of Southampton, but it was widely publicised amongst the first year students, encouraging the students to complete the application with regular prize winning draws for those who complete the application form. We acknowledge that there is no guarantee the acceptance/rejection decision would be similar to the actual decision taken if the account was a genuine one, but the students had opened such an account already and were well aware of the products and features offered by banks.

For purpose of this research, we took 21 applicant characteristics from this database. They are sex, status, number of children (Num_Children), number of credit cards (Num_Cards), income from wage, income from loan, income from contributions, other sources of income, place of education (establishment), campus, course taken (course), interests: sport, travel, music, clubbing, cinema, cars, Do-It-Yourself activities (DIY), gardening, beer and country western. The website made the students an offer of a specific account, where the overdraft limit, interest when in credit, commission-free travel money, and whether a credit card was also offered varied from student to student. The offer characteristics were chosen randomly in 25 percent of the cases and followed a decision tree which was used to ensure a variety of offers was being made in the other 75 percent of the cases. Further details can be found in Jung and Thomas (2004) and Thomas *et al.*, (2006).

After pre-processing the initial data, 305 entries were used to build the classification trees. Out of this, 214 (70 percent) was used to train the tree, while the remaining 91 entries were used to validate the trees.

pkUserNumber	Sex	Status	Num_Children	Num_Cards	Wage	Loan
1	Male	single	0	1	FALSE	TRUE
2	Male	divorced	0	3	TRUE	TRUE
6	Male	married	0	1	TRUE	FALSE
7	Female	single	0	6	TRUE	TRUE
8	Male	single	0	1	TRUE	TRUE
9	Female	other	0	0	TRUE	TRUE
10	Female	single	0	1	FALSE	TRUE
11	Female	single	0	0	FALSE	TRUE
12	Female	single	0	0	FALSE	TRUE

Table 4.1: A sample of the data from the Fantasy Student account

The data is difficult to classify well which, as one might accept when one is trying to predict a target like acceptance, especially in the case of an acceptance of a hypothetical product. We can see the results of the logistic regression classification in the next section.

4.2 Logistic Regression - Scorecard

Another method that could be used would be logistic regression. Logistic performs the same classification functions as a decision tree and indicates the significance of a variable by the value of its corresponding coefficients. The interpretation of such coefficients is as follows: the larger the value of the coefficient, the more significant the corresponding variable is to the target variable. But it is difficult to obtain and compare acceptance rate using this method, unlike in the decision tree, where the acceptance rate is displayed at the nodes of the tree. But its strength though is that it can give a prediction of the acceptance of any variant of the product offered by any applicant.

We compare the use of logistic regression on the data set (a segment of the population) as a benchmark to validate the results of the classification tree. Later in this chapter we will use it to bootstrap in order to increase the "data" set which the tree is built. The chart and table that follows show the results of a scoring exercise on a small set of data. The logistic regression and the scorecard are performed and built using SAS 9.1.3. The input variables are the 25 applicant and offer characteristics, and the target variable is the action of accepting an offer.

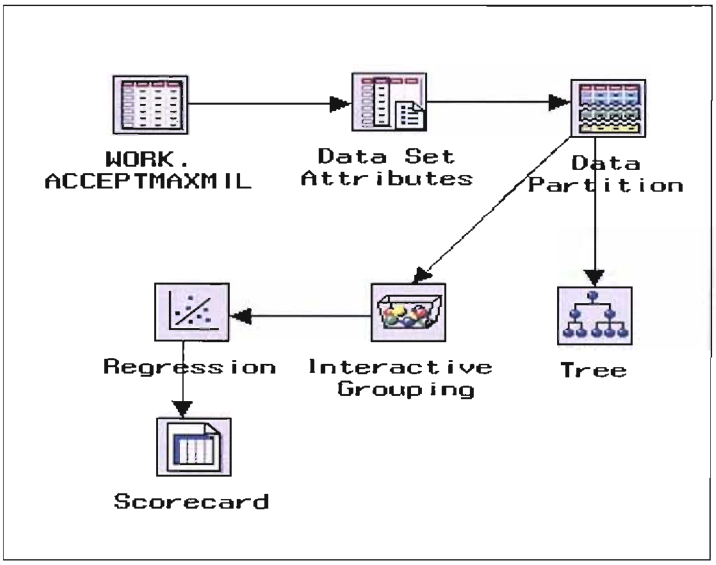

Figure 4.1: Work diagram for the building of the decision tree and scorecard

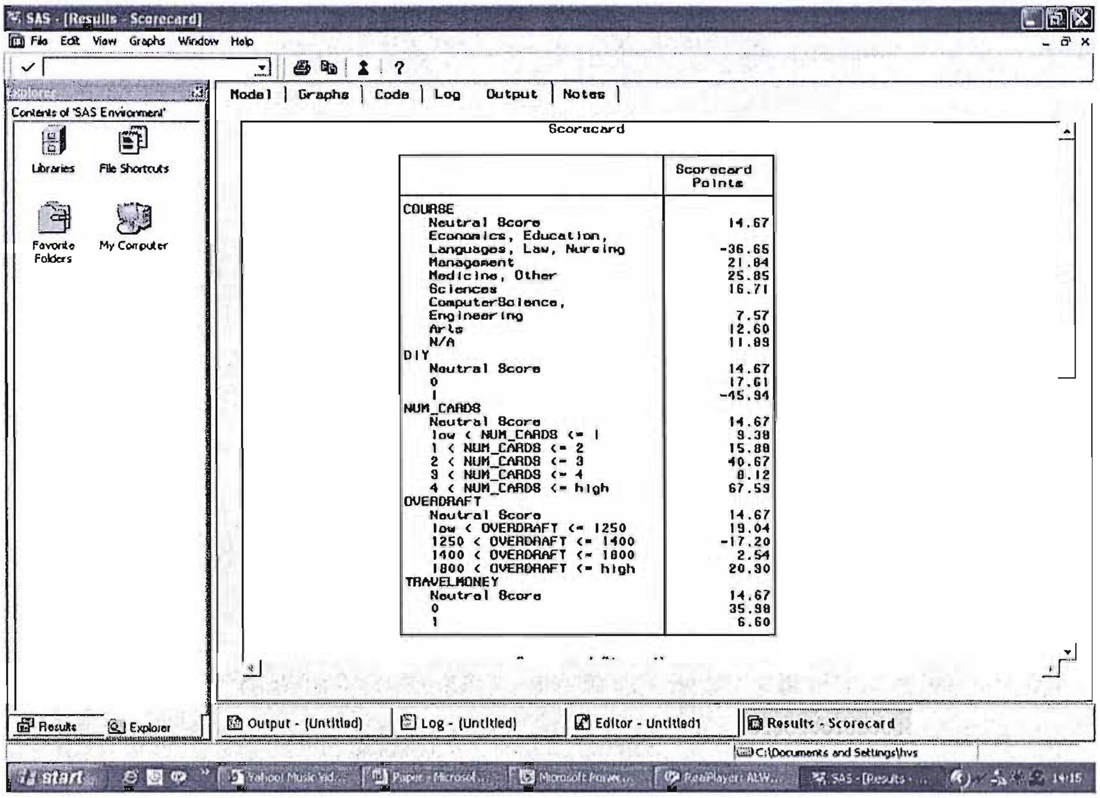

	Scorecard Points
COURSE	
Neutral Score	14.67
Economics, Education, Languages, Law, Nursing	-36.65
Management	21.84
Medicine, Other	25.85
Sciences	16.71
ComputerScience, Engineering	7.57
Arts	12.60
N/A	11.89
DIY	
Neutral Score	14.67
0	17.61
1	-45.94
NUM_CARDS	
Neutral Score	14.67
low < NUM_CARDS <= 1	9.38
1 < NUM_CARDS <= 2	15.88
2 < NUM_CARDS <= 3	40.67
3 < NUM_CARDS <= 4	8.12
4 < NUM_CARDS <= high	67.59
OVERDRAFT	
Neutral Score	14.67
low < OVERDRAFT <= 1250	19.04
1250 < OVERDRAFT <= 1400	-17.20
1400 < OVERDRAFT <= 1800	2.54
1800 < OVERDRAFT <= high	20.90
TRAVELMONEY	
Neutral Score	14.67
0	35.98
1	6.60

Figure 4.2: Output of scorecard node from Enterprise Miner 4.3 for 305 entries with 25 Applicant and Offer variables

The result of the logistic regression is used to build a scorecard as found in **Figure 4.2**. Three applicant characteristics and 2 offer characteristics are considered to have a significant effect on the acceptance of the offers made. In order to test how well the scorecard works, a Receiver Operating Characteristics Curve (ROC) is obtained. We then used the Gini Coefficient and the Kolmogorov-Smirnoff statistic to test how good the scorecard is.

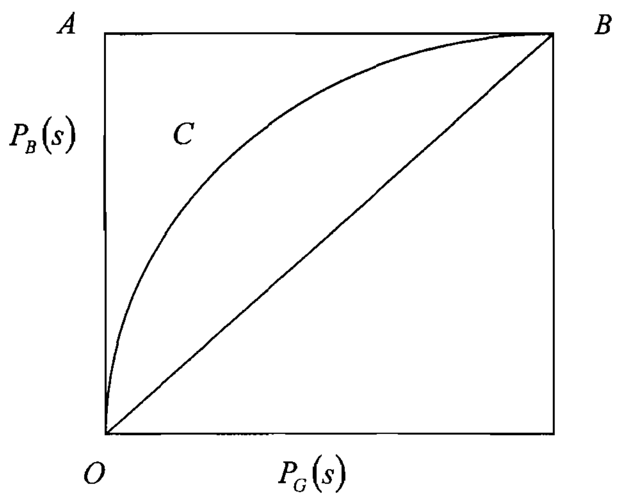

Figure 4.3: ROC curve taken from Thomas *et al.,* (2002)

The ROC curve is obtained by plotting the scores the probabilities of the bad accounts (accounts which default, ($P_B(s)$)) against the good accounts ($P_G(s)$), see **Figure 4.3**. For this case, the ROC curve we obtain is from plotting the probabilities of those who accepted the offer against those who rejected the offers. The ROC curve describes how well the scorecard classifies as the cutoff score varies. Hence the best scorecard is one which follows *OA* then *AB*. Thus at point *A*, a score s^* is assigned such that at *A*, the bads will all have a score of $< s^*$ and the goods will not. *OB* corresponds to having the same ratio of goods and bads, hence is not a good classifier.

So the further the ROC curve is from the diagonal line, the better the performance of the scorecard. This means that the larger the area between the ROC curve and the diagonal line, the better classifier the score card is. The Gini coefficient is defined as twice this area. The Kolmogorov-Smirnoff statistic is the distance between the distribution functions of the score of the good and the bad. Given that $P_G(s) = \sum_{x \leq s} p_G(x)$ and $P_B(s) = \sum_{x \leq s} p_B(x)$,

Kolmogorov-Smirnoff statistic $= \max_s | P_G(s) - P_B(s) |$.

The Gini Coefficient for this data set is 0.50729. This data set yields a Kolmogorov-Smirnoff statistic of 0.36647 and area under the ROC is 0.75364. Hence, the scorecard generated using the logistic regression is satisfactory.

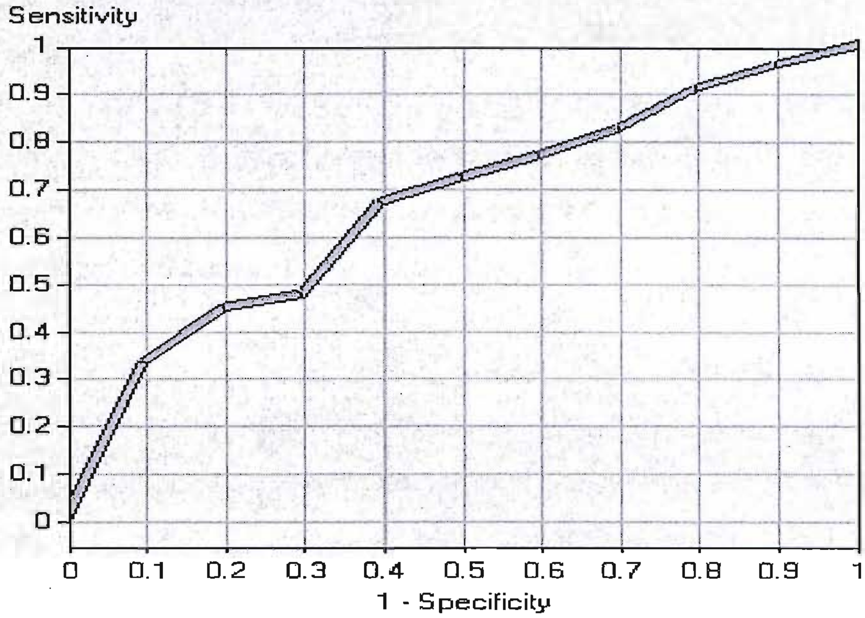

Figure 4.4: Receiver Operating Characteristics Curve (ROC) for the scorecard

4.3 Applicant and Offer Trees

The initial construction of the tree begins with using only the offer characteristics to find the offer variables for TAROT, the target variable being the acceptance of offers. This is done to determine the best offer characteristic to start building the TAROT.

The particular sequence of applicant-offer of TAROT allows questions to be formed on the applicant characteristics in the tree. The number of applicant characteristic levels in TAROT, will correspond with the number of questions the customers are asked. The following sub-sections will give details of how TAROT is built based on the data set of 305 entries with 21 applicant characteristics and 4 offer characteristics.

The acceptance model developed in the previous chapter is one that "learns" from current knowledge to predict the acceptance of the next offer. Before any prediction can be made, information that will influence the acceptance or non-acceptance of an offer has to be collected. This can be achieved by asking questions, be it in the shape of questionnaires or in real time via the internet or on the telephone. When asking for personal information, one has to be careful so as to not ask more questions than is needed for fear that the prospective customer will be uncomfortable thus reducing the probability of accepting an offer. This condition holds regardless whether this line of questioning is done face to face or through a medium like the internet.

Bear in mind that the model that we are proposing in this chapter (decision trees) does not adaptively learn (like the acceptance model in Chapter 3). The decision tree model uses past information to decide on a question (questions) to help decide which offer to make. The aim of the question (questions) would be getting "enough" information to base an offer on but asking the "right" questions. The difficulty of such a situation is choosing the "right" questions to ask. The "right" questions here are defined as a minimum number of questions with responses that will give enough data about the potential customer so that an "acceptable" offer made to him, or her; which is one with a high probability of acceptance. In such a situation, it is crucial to learn as much as possible by asking as few questions as one can. If each experiment to get information through asking incurs the same cost, hence minimising the expected number of questions will also minimise the costs involved (Raghu *et al.*, 2001)

The unique sequence of building for TAROT requires the flexibility of interactive training found in the "Tree" node of Enterprise Miner 4.3. The interactive version allows for the analysis of the effects of interaction between the applicant variables and the resulting offer characteristic labelled to the corresponding leaf node. It also allows for selection of (a limited choice) variables to be implemented within the tree in real time. Hence the resulting interaction with the acceptance rate can be analysed in real time and with relative ease. However, if a certain variable is not significant enough to be considered at a particular split of the tree, the interaction with the acceptance rate can be calculated offline. This is done manually and is time consuming.

4.3.1 Offer Characteristics

The offer characteristics will play an important role in building TAROT and guide in making the decision of which offer to extend. The selection will be based on the offer that has a high acceptance. The process will depend on the given responses of the customer to the question/questions asked. This is effectively matching the offer to the responses. The decision tree that follows show the best tree derived from only the offer characteristics. It shows the best/most popular offer for that sample. The more intense the colour and boldness of the lines, the more popular that variable is with regards to the acceptance of offers.

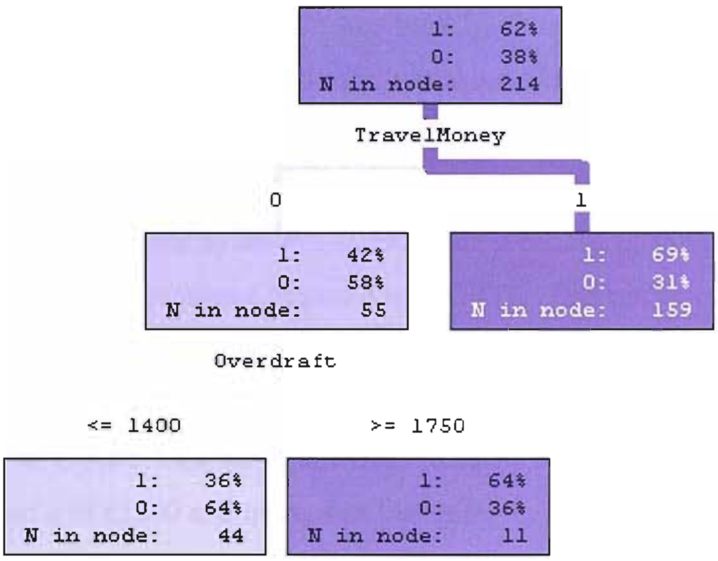

Figure 4.5: Offer characteristics only classification tree

Notice again that this is a classification exercise showing that of 62 percent of the total 214 people have accepted offers. 69 percent of the 159 people that had travel money offered to them accepted while 42 percent of 55 people who were not offered travel money accepted offers. But from this 55, 44 were offer an overdraft of £1400 and less, with 36 percent accepting while the others were offered an overdraft of £1750 and above, 64 percent of those accepted.

This decision tree only gives information on the most popular offer characteristics in the sample data. So, the next step would be to build a decision tree which has all the applicant and offer characteristics and look into the acceptance rate on this tree. But again, one will realise that such an exercise will produce a classification tree with no means of helping decide of which offer to extend to which kind of customers. This tree is included in this chapter to serve as a reference to compare with TAROT.

4.3.2 Classification Tree Consisting of Both Applicant and Offer Characteristics

Figure 4.6 is the classification tree built by Enterprise Miner 4.3 which is the best tree using all 25 applicant and offer variables. Notice that there is a lack of coherent sequence and understanding that allows for interpretation of the classification to result in the identification of an offer of any sort, let alone an offer with a predicted high acceptance rate by the customer. This is caused by the lack of order in the tree compared to the sequence imposed in TAROT. There is no clear way of identifying an appropriate offer due to the lack of a sense of logic as the variables are randomly scattered among the nodes of the tree. For example, a classification from the tree looks at people who have not been offered travel money but have been given overdrafts of £1500 and more, and like going to the cinema.

Another classification is among those who have accepted an offer with commission-free travel money. We look at those who have 2 credit cards at their disposal. If they are studying Arts and dislikes DIY but enjoy travel, then they should be offered an overdraft of more than £1125. However, this is circuitous. If they accepted the offer with commission-free travel money, was it for less or more than £1125. For this same group of people what should we do if they are not offered commission-free travel money? This is an example showing how difficult is it to formulate questions from the applicant characteristics of the tree and more importantly, how hard it is to identify a good decision.

In order to make the tree usable, we need to at the top split on applicant characteristics which segment the population into groups. Then we can have splits on the offer characteristics to identify which offer to make to what segment. The difficulty in using a classification tree in the normal way, even with the first level split being forced to be an application characteristic is that these splits are made without considering that there may be a bias in the accept/reject decision because different offers were made to different applicants. So we need to restrict this bias in some way even if we cannot remove it entirely. Thus we try to identify what are the most important features of the offer in terms of the accept/reject decision and look at what happens if these types of offers are made. This will restrict the number in the sample who will be considered at a step but the usual policy of defining a node as an end node of the tree if there are insufficient cases should prevent analysis being done on too small samples. Then one can use the iterative methodology in **subsection 4.3.3** to build the TAROT tree.

The foundation of the building a TAROT of i levels of the tree is that for the first m levels, the competing splits will consist of only applicant variable based splits which the final, lowest level, n is the offer to be made. The strong classification powers of the CART will ensure that the variables that are being considered are significant.

So, it is obvious that at certain stages of building TAROT, some variables will have to be "forced in" the tree. What is certain is that each variable that is chosen is of significance even if at a previous node, the variable chosen was "forced in". This is also ascertained by the Tree Node in Enterprise Miner 4.3 in SAS 9.1. Hence there is a general assurance that although this is not an optimal tree, it is the best possible tree built based on the restrictive nature of the order of variables and taking into account the variable chosen at a prior node in a level. We do not assume the tree gives an optimal segmentation and subsequent offer in any sense but it is a good compromise between finding "locally optimal" solutions and offsetting the bias that occurs in such "local optimal" solutions which do not recognise that the data consists of different offers to the different customers. These trees are presented in the next section.

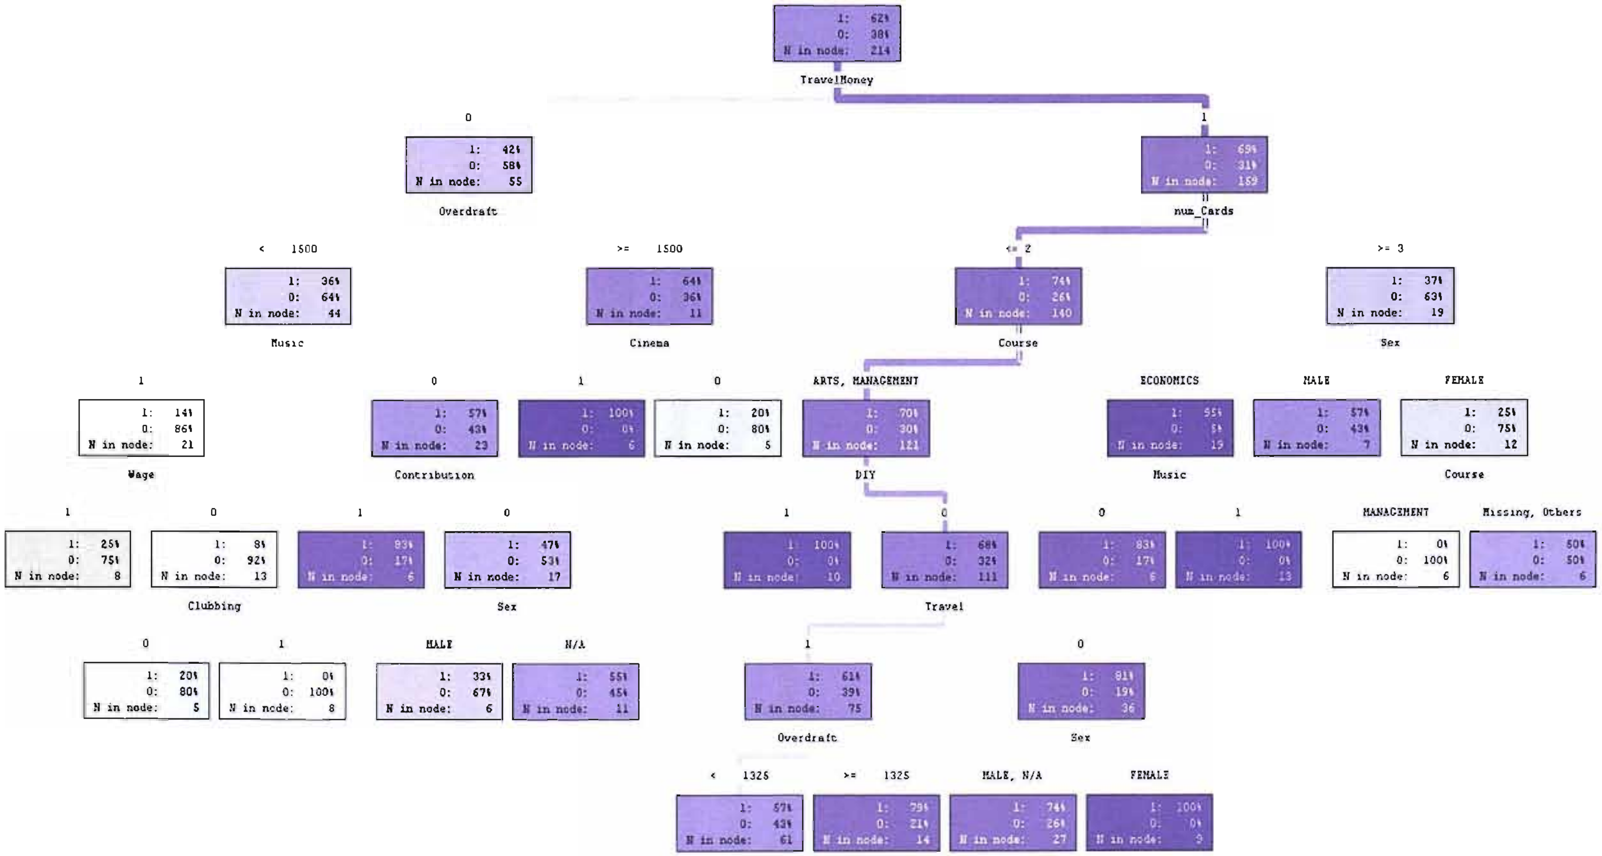

Figure 4.6: Applicant and offer characteristics classification tree (with no "initial classification using applicant characteristic" rule)

4.3.3 Building TAROT

TAROT can be built to accommodate m number of questions, m = minimum number of questions to be asked. In this section, a step-by-step guide on how TAROT is built has been given. There are instructions on how to build a one-question, one-offer (1, 1)-TAROT, and a m-questions, one-offer (m, 1)-TAROT.

A one-question TAROT is by far the easiest among the TAROTs to construct. The following are the instructions. First identify the best offer from an only-offer-characteristic tree. By keeping this offer characteristic in the first level of the TAROT, find the applicant variable that gives the best acceptance rate for that offer. Then, replace the offer characteristic in the first level with the chosen applicant variable at the first level, and experiment with the rest of the offer characteristics and chose one that has the highest acceptance rate. Once the offer is chosen, the one-question TAROT is complete.

A (m, n)-TAROT assumes that there is a maximum of m levels of applicant characteristics and n levels of offer characteristics in the tree. Such a tree would lead to an offer process that uses at most m questions in order to segment the population and there are at most n features of the product that will be varied from the standard as part of the offer. It is necessary to define a standard value for each of the features of the offer (usually the least expensive to the lender) and it is assumed this is what is offered if that offer characteristic does not occur in the tree.

For the construction of a m-question, n-offer TAROT, the following steps were taken:

1. Build an "up to n-level offer characteristic only tree" on accept/reject decision and from this one can identify the combination of features that give the highest acceptance rate and which has been offered to sufficient applicants.
2. Use this highest offer combination at the top of a tree and add a last level of the best split of the application characteristics.

3. Now take this best "bottom level" application characteristic split and force it in at the top of a classification tree and allow up to n-level of lower level offer characteristics splits.
4. Take the branch of the n-level offer characteristics splits that lead to the highest acceptance rate and is applied to a sufficiently large proportion of the population – one branch from each of the subsets arising from the application characteristic splits – and add a bottom level application characteristic split. The choice of what percentage of the population "is sufficiently large" is a parameter of the methodology.
5. Force in these "bottom level" application splits below the current application levels in the tree and then allow up to n-levels of offer characteristic splits.
6. Repeat steps 4 and 5 until the m-levels of application characteristics have been introduced or the characteristic split in 4 is not significant.

When the tree is completed the top level application splits give the segmentation of the application population, while the lower level offer splits describe the different changes that have to be made to the standard offer for each of these segments. The idea of the tree is that one is concentrating at each stage on the offers with the highest acceptance rates when deciding which further segmentation of the application characteristic to make.

The following is a tree built for the (2, 1)-TAROT. Note the best offer from the tree built from offer only characteristics (**Figure 4.5**) is Travel Money. Hence we find the best first applicant variable with the highest acceptance rate for Travel money is Number of Cards denoted as "Num_cards" in the TAROT.

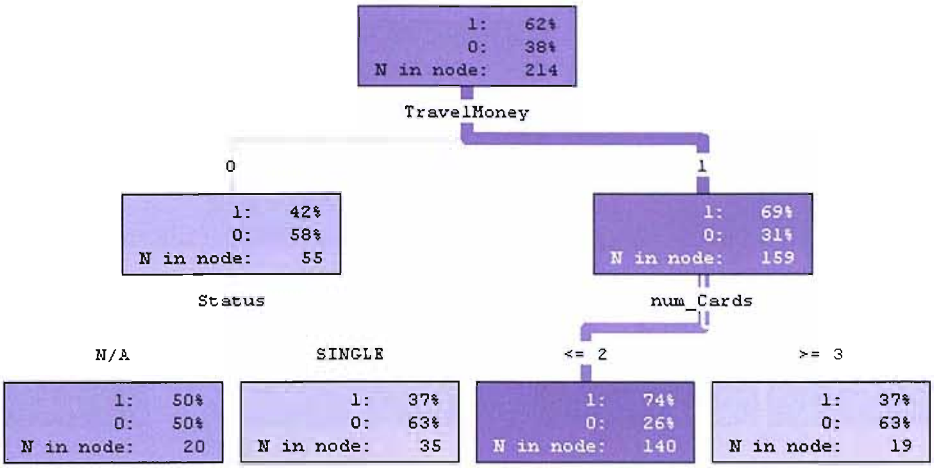

Figure 4.7: Finding the first applicant variable split from the best offer taken from the offer characteristics only tree

Then, the applicant variable is then placed at the first level, and we find the best offer to suit the groups resulting from the split.

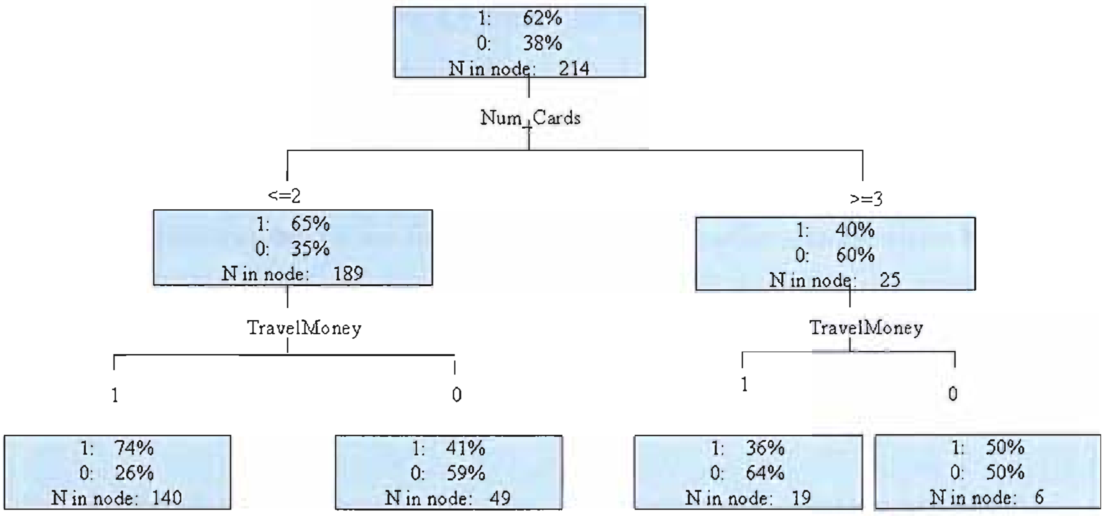

Figure 4.8: Finding the best offer for the applicant variable identified

The next step is to find the best applicant splits for the branch of giving Travel Money for customers who have 2 or less cards, and for the rest, the best applicant split for not offering Travel Money.

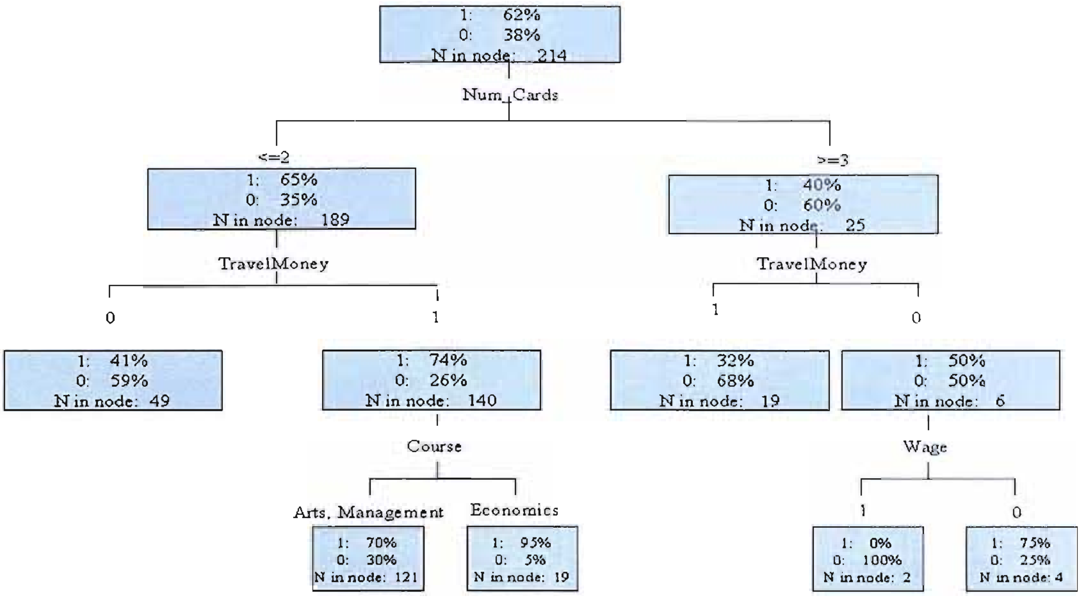

Figure 4.9: Finding the second applicant split from the branches of the offers with the higher acceptance probability

By following Step 4, we get **Figure 4.9**. Note that the best applicant characteristic for those with 2 or less cards, that have been offered Travel Money, is Course. For the customers with 3 and more cards, the best applicant characteristic from the branch of not offering Travel Money is Wage. Hence by placing these second applicant characteristics below the first, we then allow any offer characteristic to be placed under the second applicant characteristic and choose the offers that have the highest acceptance. These are illustrated in **Figures 4.10** (Interest), **Figure 4.11** (Travel Money), **Figure 4.12** (Credit Card) and **Figure 4.13** (Overdraft).

We then compared the acceptance rate of the offer characteristic and the best offer with the highest acceptance rate is chosen as in the TAROT in **Figure 4.14**. An example would be the offer characteristic for customers having 2 or less credit cards who study economics. The strongest acceptance rate among is between Travel Money with 95 percent acceptance for a sample of 19 people. Although both Interest and Credit Cards have additional 2 and 1 people in them, that is not significant enough when compared with their acceptance rate of 90 percent. Hence Travel Money is chosen.

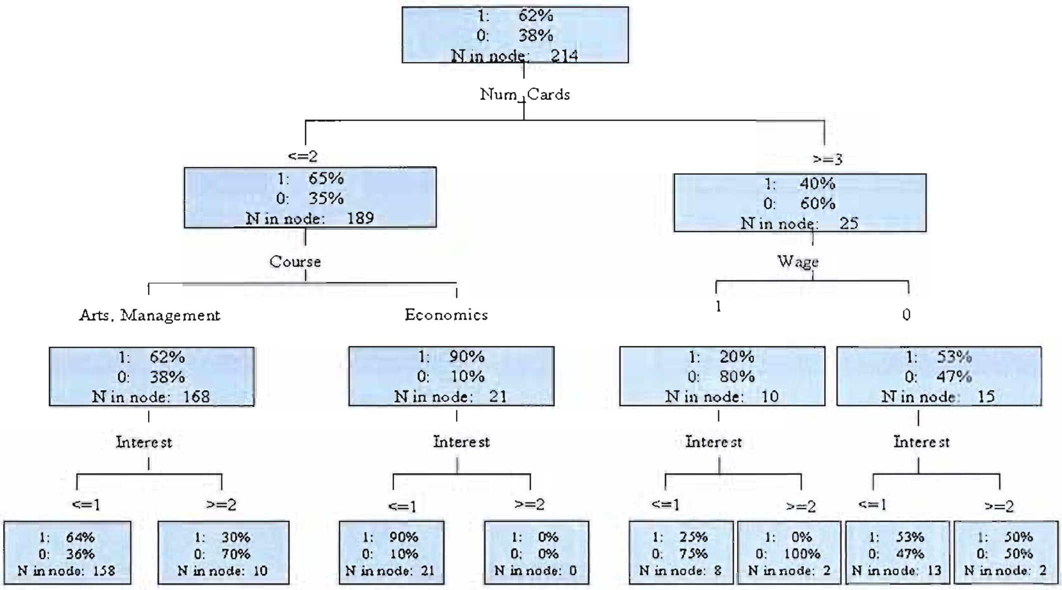

Figure 4.10: Applicant and offer characteristic tree- Interest

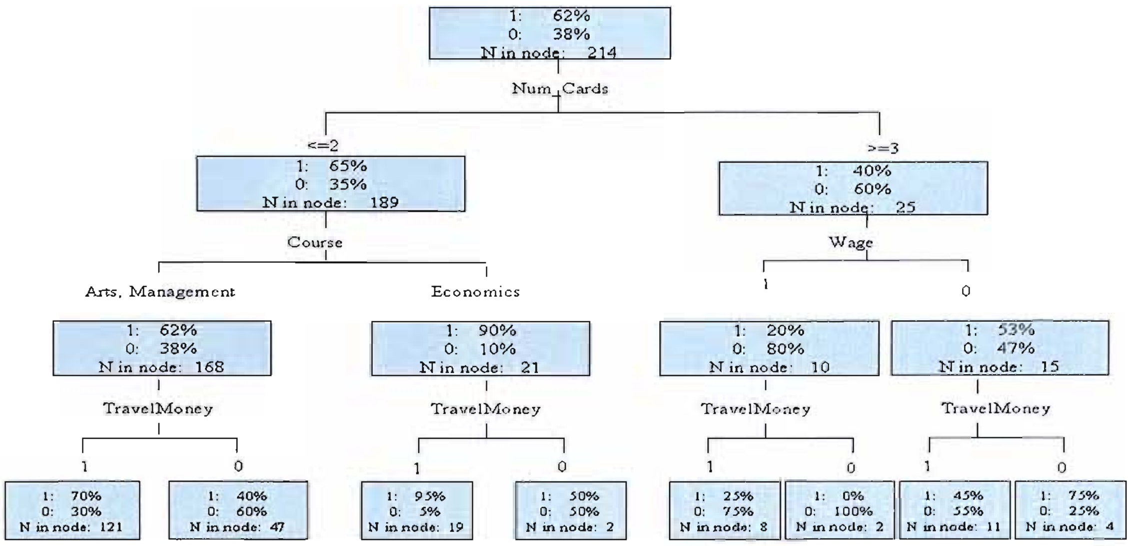

Figure 4.11: Applicant and offer characteristic tree- Travel money

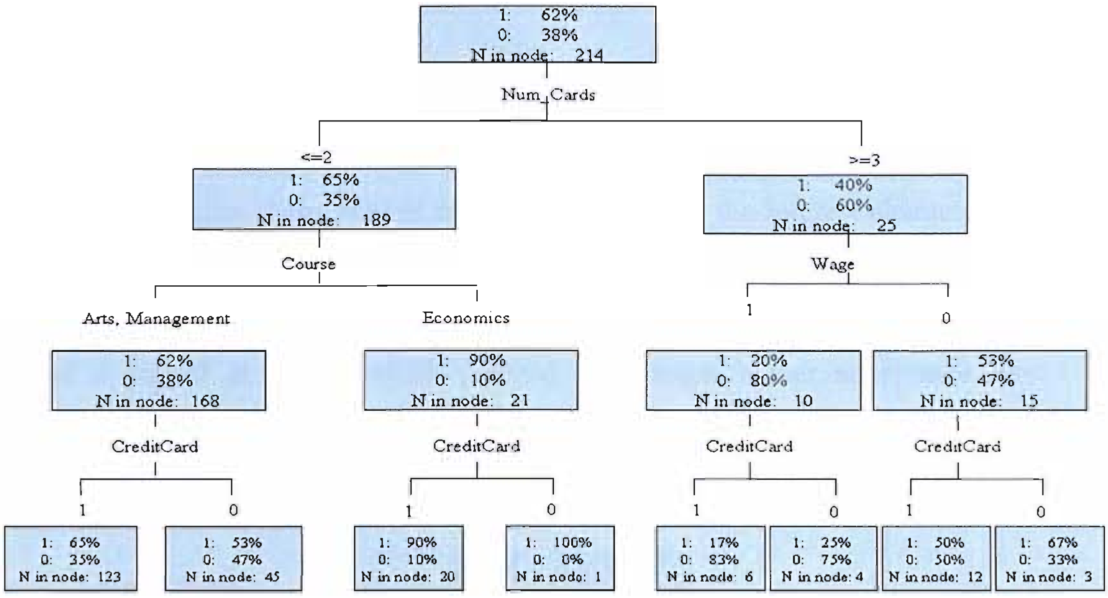

Figure 4.12: Applicant and offer characteristic tree- CreditCard

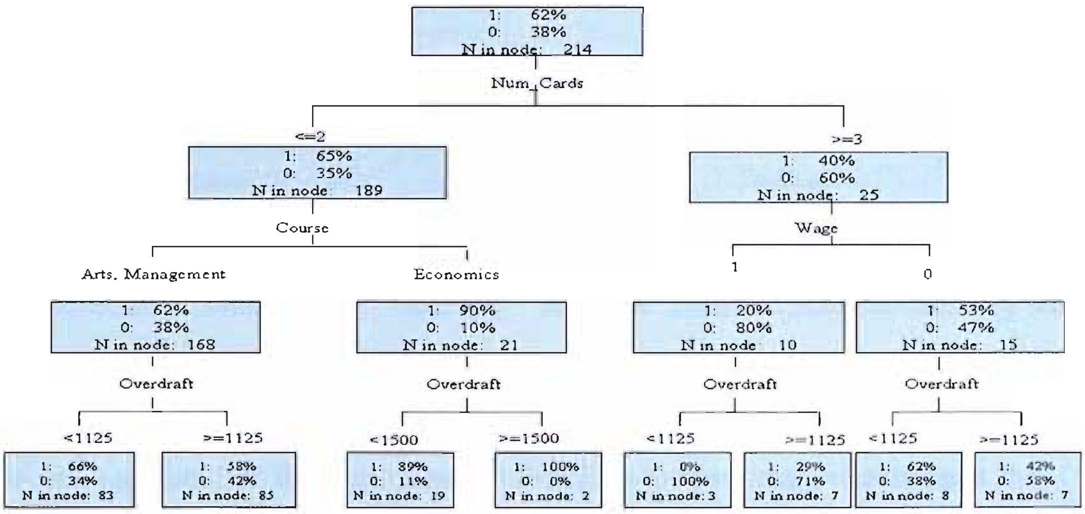

Figure 4.13: Applicant and offer characteristic tree- Overdraft

In **Figure 4.11**, Travel Money was the most significant offer for Arts and Management. Hence it is maintained for that course variable. With the best second applicant variable found, the identification of the best offer commences. The following trees show the acceptance rate given Travel money, Interest and Overdraft as the offer extended.

The overdraft limit for the wage split is £1125 and not £1500 as for the Economics split. This is because at the cut-off of £1500 will lead to the Interest characteristic giving the highest acceptance rate. But looking at the following one level split decision tree, Overdraft is more significant than Interest. So Enterprise Miner allowed a cut-off at £1125 which resulted in a much higher acceptance rate compared to Interest at the 2-level decision tree. This process resulted in the following tree for a 2-question TAROT for this data set.

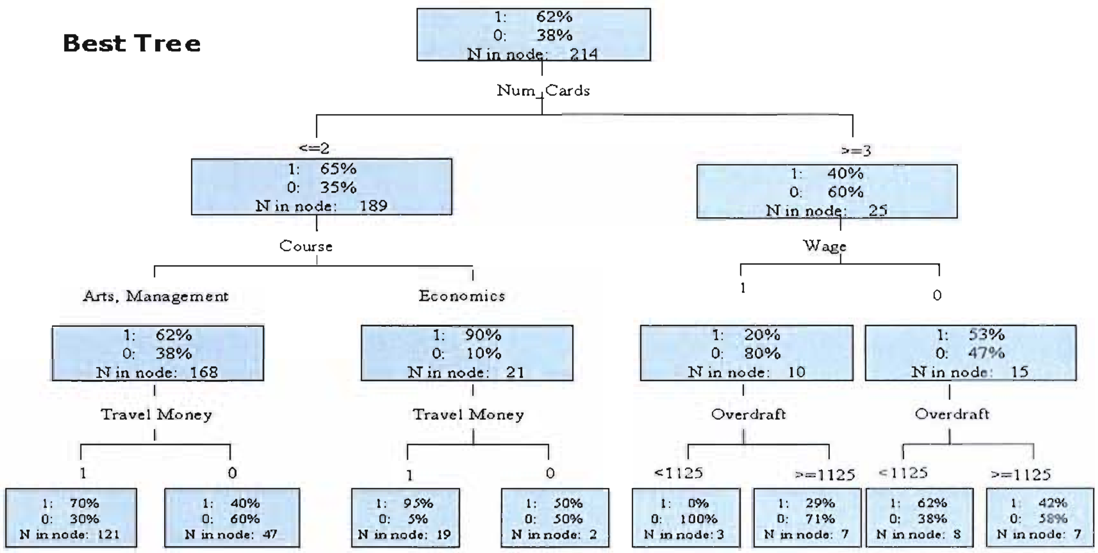

Figure 4.14: (2, 1)-TAROT for the Fantasy Student Account

The tree suggests that the first question should be on how many credit cards the applicant has. Hence, the responses split the population into those with 0, 1 and 2 cards and those with 3 or more credit cards. The second question for the first group is that subject they study and for the second group, whether they receive a wage or not (i.e. do they have a part-time job). Notice the questions asked are adaptive in that they depend on answers to the previous question.

The problem with this tree is that we have forced the offer decision to be the same for all groups – to whom do we give travel money? The tree suggests we give travel money to those with 2 or less credit cards. It is assumed that if an offer feature is not significant in the tree, we will leave it at some pre-assigned level.

The obvious question is whether it is better to vary the offer on other characteristics to the few applicant groups we have identified. So **Figure 4.10** to **Figure 4.13** are the result of allowing this third level splits to be on any offer characteristic. In this case, we see that it is both travel money and overdraft which most affect the acceptance rate. In this case, it is best to offer commission-free currency exchange to those with 2 or less credit cards and to offer overdrafts of more than £1125 to those with more than 2 credit cards but have a wage, but overdraft of less than £1125 for those who have more than 2 credit cards but no wage.

Notice that the waged, more than 2 credit card groups has an acceptance rate of less than 30 percent and so one might decide not to make them any sort of offer.

So, what can be seen from the optimal TAROT tree is the following:

a) If a customer possesses two or less credit cards and is a student of Arts or Management, hence the best offer to make to this individual is commission free travel money.
b) If a customer possesses two or less credit cards and is a student of Economics, hence the best offer to make to this individual is commission free travel money.
c) If a customer possesses three or more credit cards and is earning a wage, the best offer to make to this individual is overdraft of £1125.
d) If a customer possessed three or more credit cards and is not earning, the TAROT predicts the best offer to make is an overdraft of less than £1125.

There also exists an alternative best tree as presented in the next section:

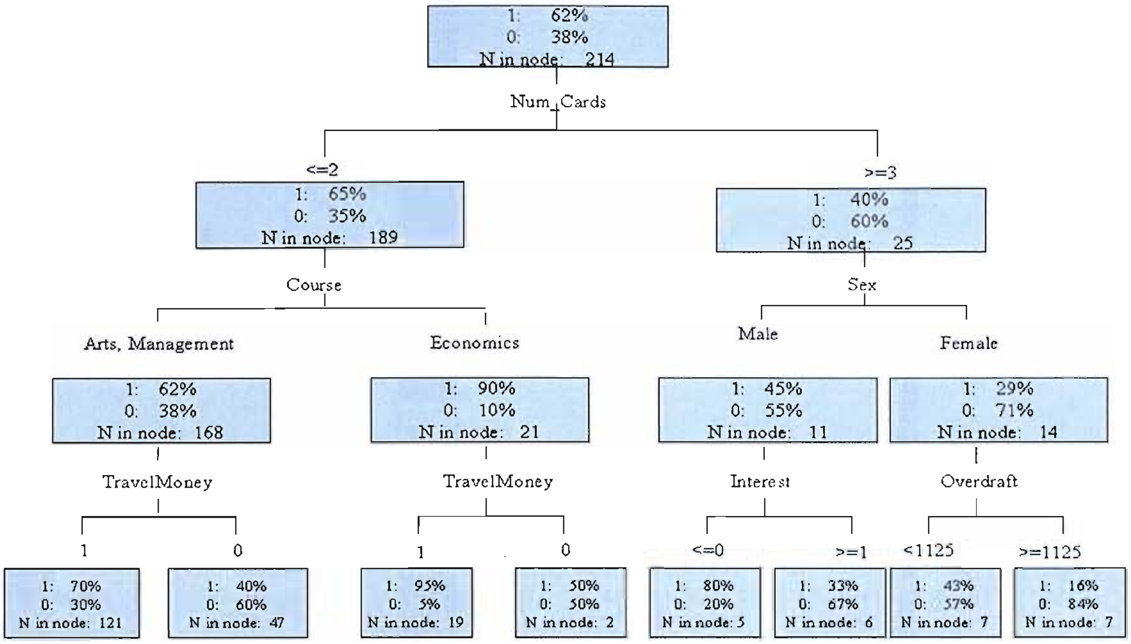

Figure 4.15: Alternative TAROT for the data set for the Fantasy Student Account

There is an alternate which is quite similar to the tree in **Figure 4.14**. In the first tree, we were only allowing offer or applicant characteristics to vary in a tree. Suppose we allowed the final level of applicant characteristic and the offer characteristic both to vary. We get this tree in **Figure 4.15**, which leads to better acceptance rates. In this tree, we check the set of those with 3 or more credit cards and if male, give them 0 percent interest on credit, and if female an overdraft of less than £1125. These seem odd results in that in both cases, one is making an apparently worse offer. This could be due to the data coming from a fantasy situation, but it is recognised that some students do not like having large overdraft limits, while others believe that interest on current accounts is compensated for by hidden charges and is of no benefit if you are never in the black.

Notice that if we take the standard classification tree (**Figure 4.6**), it does not split the population into segments each which results in one offer. So, it is not appropriate for identifying what offer to make.

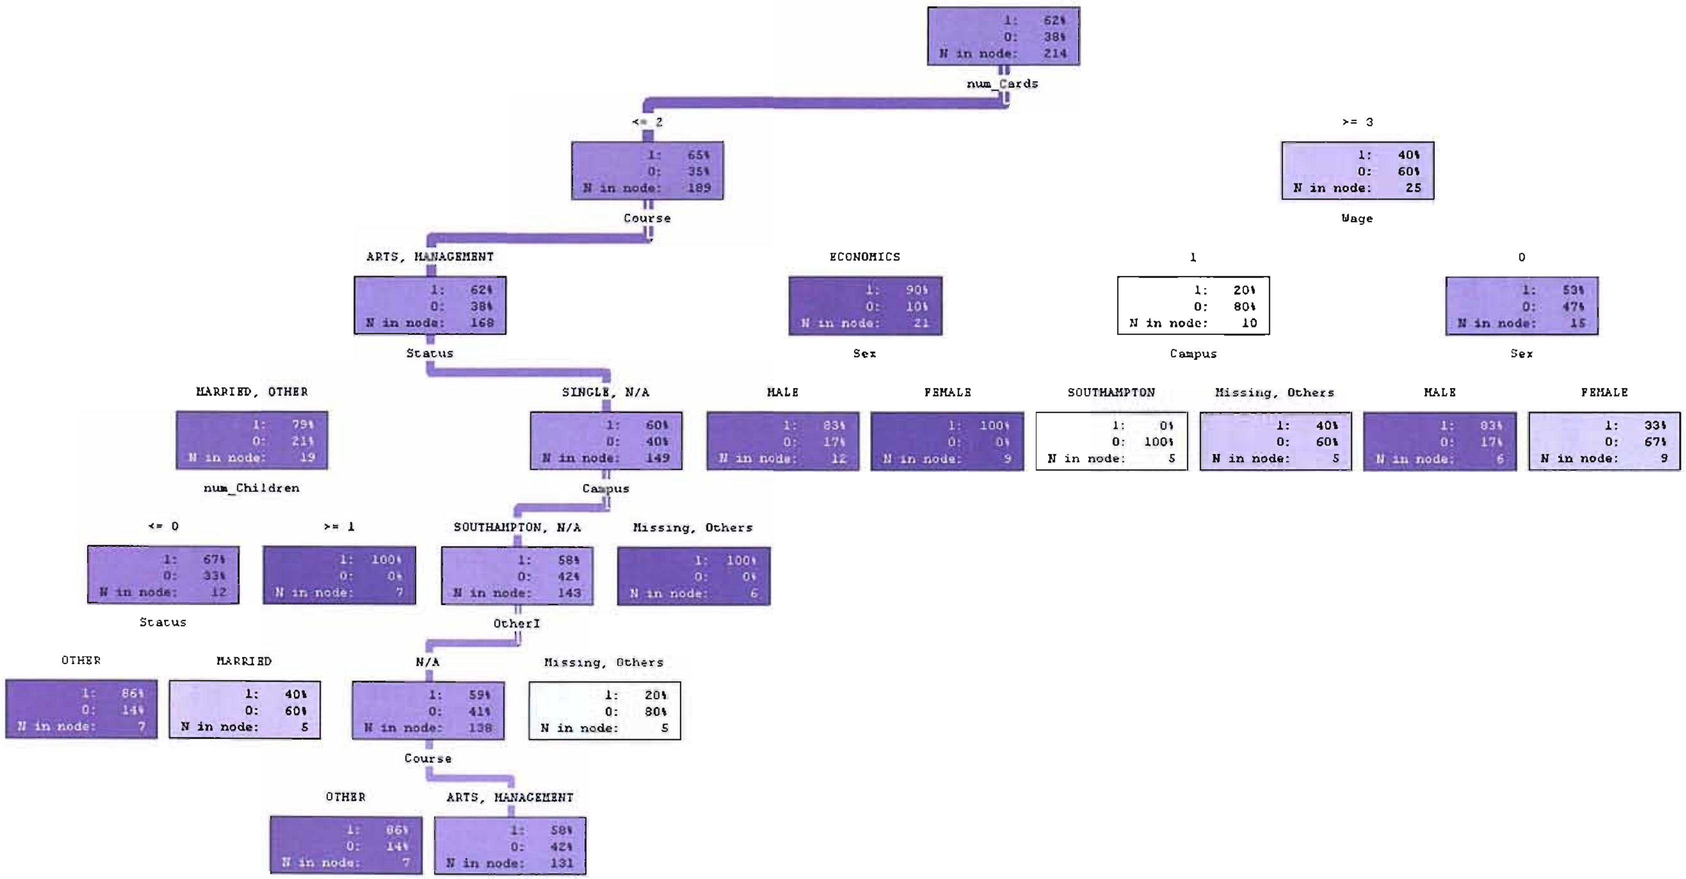

Figure 4.16: Best classification tree for Student Fantasy Account data with most significant applicant characteristic at root node

4.4 Decision Tree - Probabilities of Acceptance

Small data set- a way of dealing with this

With 305 entries, there will be concern with the small data set. In an attempt to deal with this, hypothetical acceptances are generated to gather a much larger data set than that of the actual acceptances. This is a form of "bootstrapping" which may make TAROT more robust.

Firstly, we take the 5 variables of the scorecard in **Figure 4.2**. Then, we look at the Weights of Evidence and the point estimates from the logistic regression as found in **Figure 4.17**.

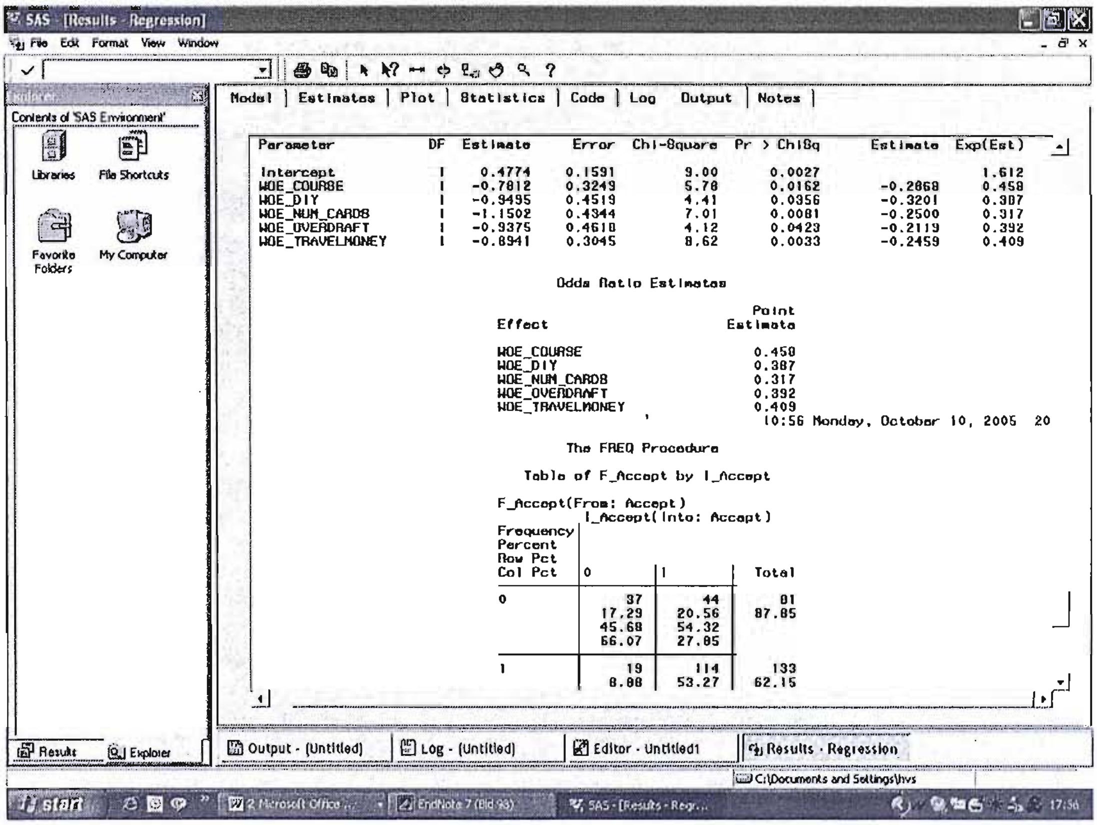

Figure 4.17: Results of logistic regression of Student Fantasy Account

4.5 Calculation of the Probabilities of Acceptance

Let

S = Score of variable

p = Probability of acceptance of offer

c_0 = intercept

c_n = value of weight of evidence for a field accompanying applicant variable n

d_n = value of weight of evidence for a field accompanying offer variable n

A_n = Attribute n

TM_1 = accept offer of travel money

TM_0 = reject offer of travel money

Ov_1 =overdraft $\leq$ £1250

Ov_2 =1250 $<$ overdraft $\leq$ 1400

Ov_3 =1400 $<$ overdraft $\leq$ 1800

Ov_4 =overdraft $>$ 1800

$$S = \log\left(\frac{p}{1-p}\right) = c_0 + c_1A_1 + c_2A_2 + c_3A_3 + d_1TM_1(\text{or } + d_2TM_0) + d_3Ov_1$$

$$(\text{or } + d_4Ov_2 \text{ or } + d_5Ov_3 \text{ or } + d_5Ov_4)$$

Having built the logistic regression, we can use it to increase our data set 8-fold. Each existing entry in the data corresponds to one specific offer (one of four overdraft levels and whether or not to offer commission-free money exchange). We can use the scorecard to obtain the "scores" for the likelihood of that applicant accepting the other 7 combinations of offers. We use the logistic transformation between scorecard probability to obtain the equivalent probability and build the tree using these as the target variable values. The transformation is as follows:

$\frac{p}{1-p} = e^S$, hence $p = \frac{1}{1+e^{-s}}$.

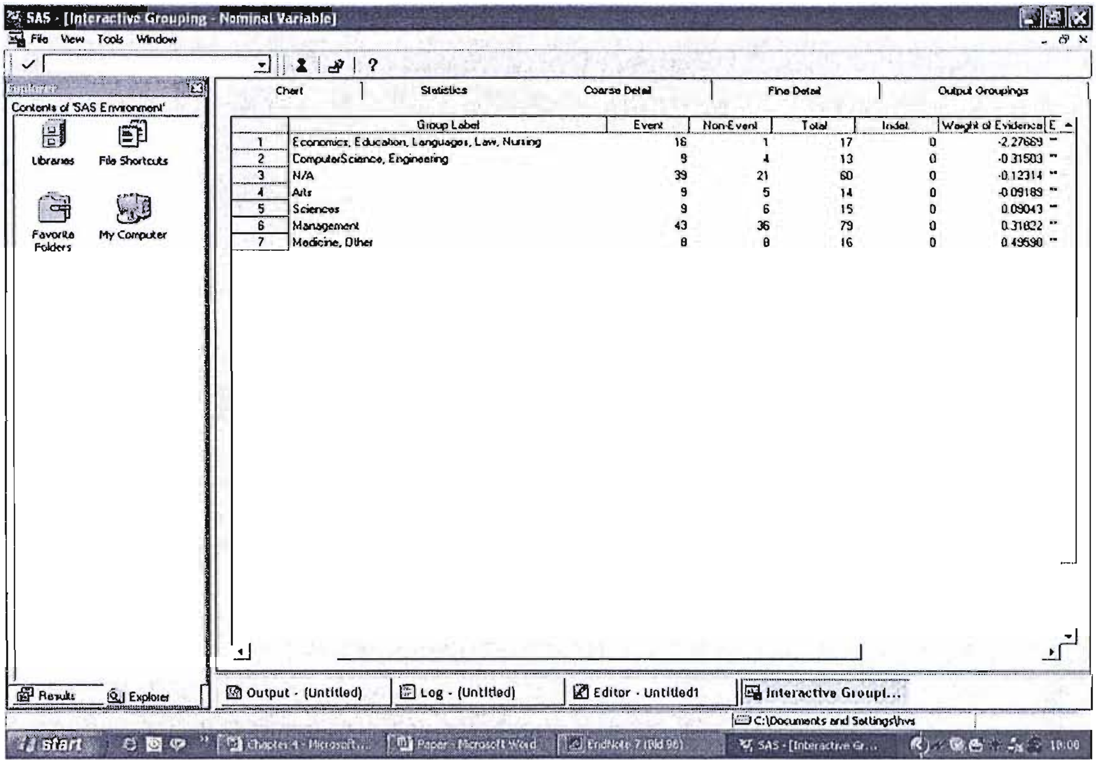

Figure 4.18: Weights of evidence of fields of the course variable

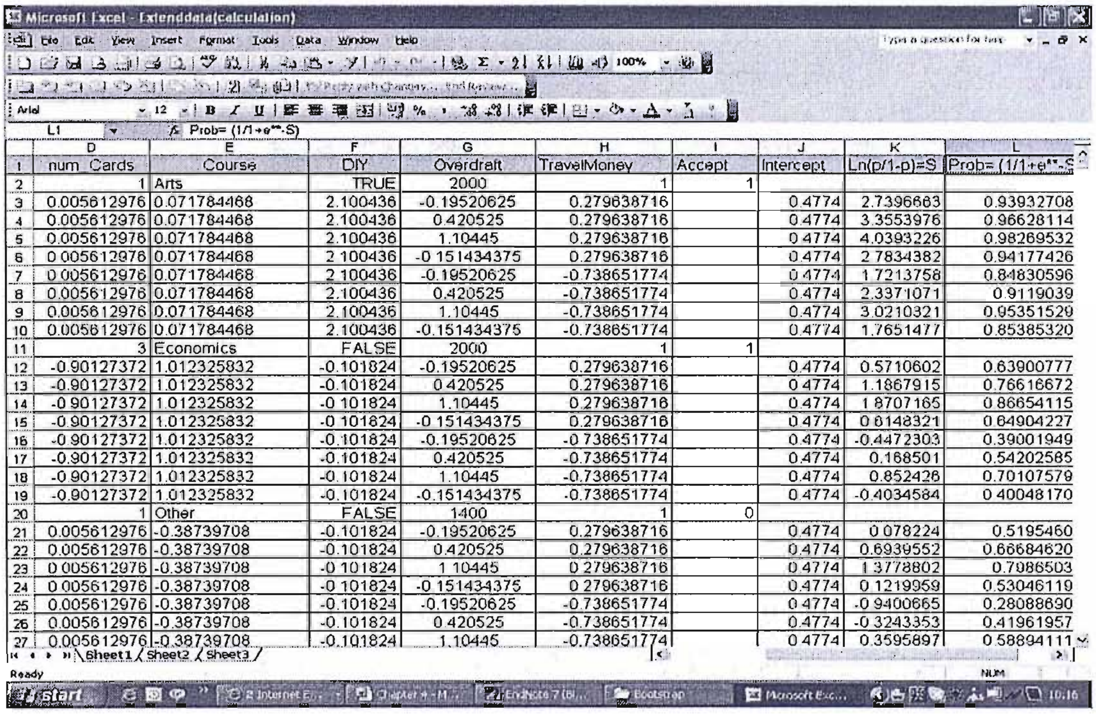

Figure 4.19: Example of calculation of hypothetical acceptances for Fantasy Student Account

The following (2, 1)-TAROT was built from the extended data:

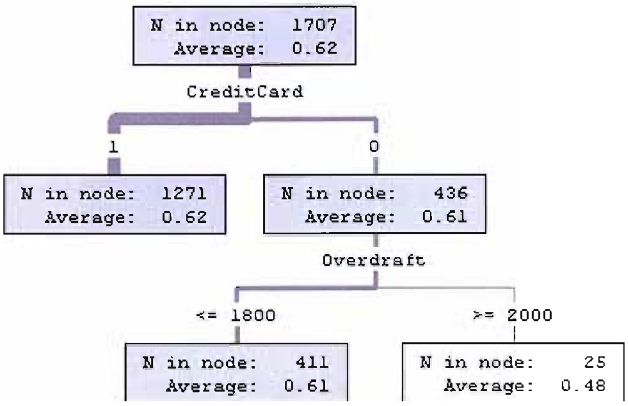

Figure 4.20: Offer characteristics only classification tree

Note that the offer characteristic selected is CreditCard. We use this offer to find the first applicant characteristic.

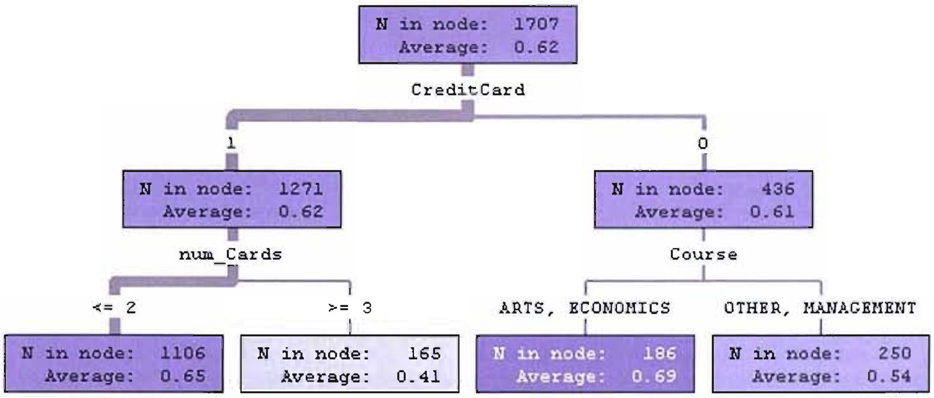

Figure 4.21: Finding the first applicant characteristic split

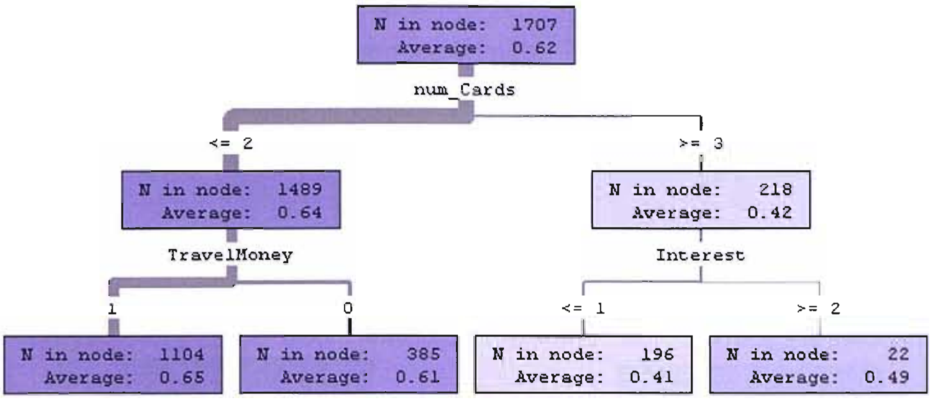

Figure 4.22: (1, 1)-TAROT from extended data

Then, using the methodology for building the (2, 1)-TAROT, we utilise the offers from the (1, 1)-TAROT to find the second applicant split from the offer branch with the highest acceptance (**Figure 4.22**).

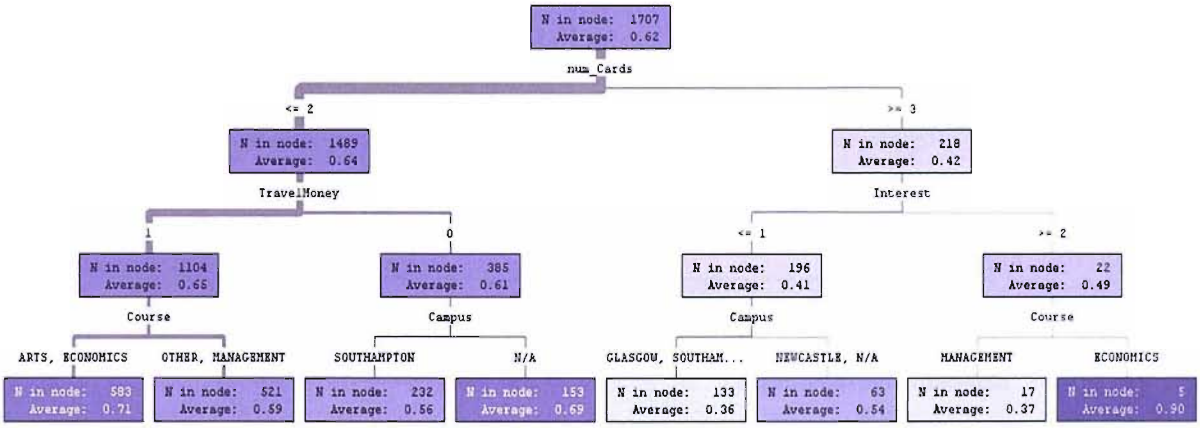

Figure 4.23: Finding the second applicant split from the offer branches

Course and Campus are the selected applicant splits. Looking at **Figure 4.23**, based on the methodology, the branch of 2 or more percentage of interest should have been chosen. Hence the chosen applicant characteristic for that branch should have also been Course. But the leaf with the highest average rate of acceptance (0.9) has 5 entries; hence it is too small a sample to be considered. The next leaf has an acceptance average of 0.37, which is lower than the acceptance average of 0.54 from the Newcastle, N/A set from the Campus branch. Hence the Campus variable is chosen as a split for people with 3 or more credit cards. (see **Figure 4.24**).

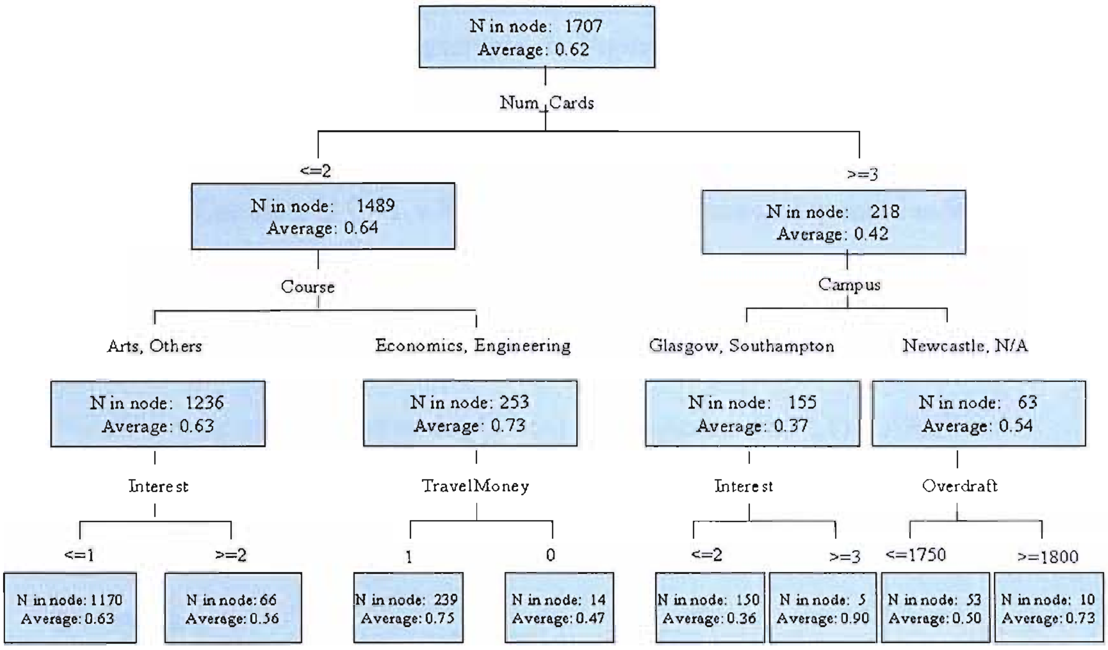

Figure 4.24: (2, 1)-TAROT from extended data

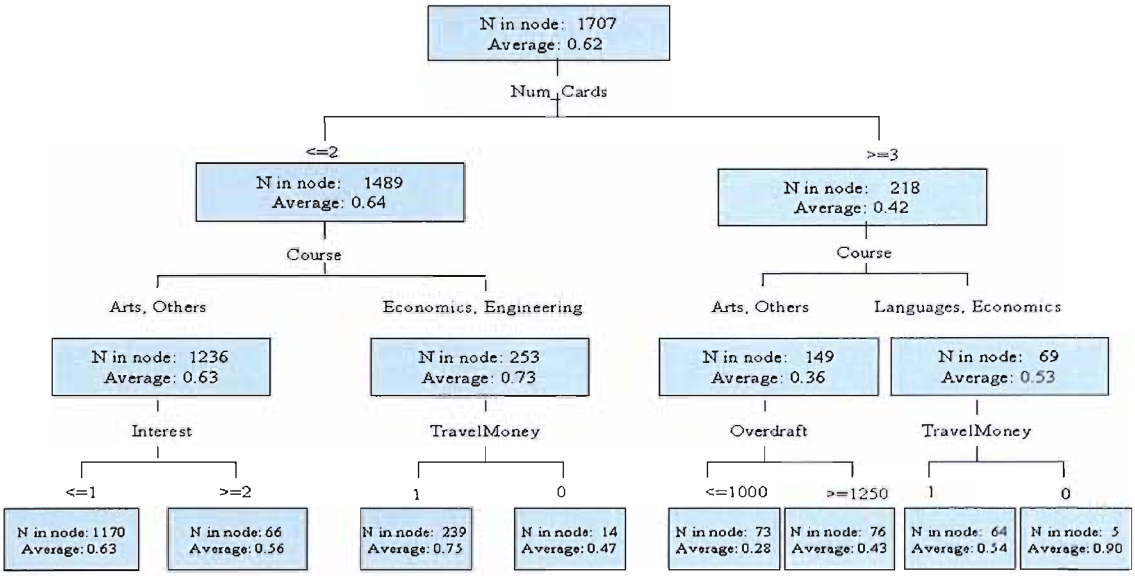

Figure 4.25: Alternative (2, 1)-TAROT from extended data

One will come across various alternatives for the (2, 1)-TAROT tree which have some different applicant split. In some cases, there is not a large difference in the acceptance average. Then, one will have to consider the sample size of the node so as to choose the best alternative. For example, in **Figure 4.25**, on the right branch, Course is the alternate applicant split for people with 3 or more credit cards. By looking at the right branch of **Figure 4.23**, the acceptance probabilities for which lead to the variable Campus, is 0.41 while the acceptance probability that leads to Course is 0.49. But the sample of the population taken leading to the Course split is relatively smaller than the sample taken for Campus. So, we built two TAROTS (**Figure 4.24** and **Figure 4.25**) and compared the acceptance probabilities. Based on this, we chose Campus to be the second applicant split. Hence the (2,1)-TAROT in **Figure 4.24** is the best.

4.6 Conclusion

The decision trees proposed do help in deciding offers, with high acceptance rates. We have developed the trees on a fairly small data set and so it seems unrealistic to go beyond a two questions, one offer characteristic set of decision. However, the approach should work whatever the m, n in the (m, n) TAROT.

Chapter 5

Choosing Questions

5.0 Introduction

We introduced the Acceptance Model in Chapter 3. It is a model based on dynamic programming with elements of Bayesian learning to predict the next best offer to extend the customer. Then in Chapter 4, we wanted to know whether it is possible to select an offer with a high acceptance rate by asking a question or two. We developed TAROT and proceeded to show how TAROT is used to aid the selection of an offer to extend to a customer. In this chapter, we will use the acceptance model to predict the next offer to extend to the customer based on the response given to a number of questions asked and which questions should be asked.

Following Chapter 3, we will consider only two different offers even though as in Chapter 3, our results can be extended to more offers. These two offers are Offer *a* and Offer *A* where $p_a > p_A$ and $g_A > g_a$ where p_i is the probability of accepting offer i and g_i = profit from offer i where $i = a, A$.

Recall the basic problem from Chapter 3. We retain throughout the entire chapter the assumption that the number of potential customers has a geometric distribution with a parameter β, so the chance of the current customer being the final one is $1-\beta$. In addition, we assume that if a customer rejects variant i, then he will reject any other worst variants, j where $j > i$. Similarly if he accepts variant i, he will accept all better variants, j when $j < i$. For ease of mind, instead of using 1, and 2 to describe the offers, we will use *a* and *A*. Hence if a customer rejects Offer *a*, he will also reject Offer *A*. If he accepts Offer *A*, he will also have accepted Offer *a*. However, bear in mind accepting Offer *a* does not compel the customer to accept Offer *A*. We ensure this by defining a set of conditional probabilities:

p_a = Probability of accepting offer a

p_A = Probability of accepting offer A

g_a = Profit from offer a

g_A = Profit from offer A

h_A = Probability $\{\text{accept Offer } A \mid \text{accept Offer } a\}$

So $p_A = h_A p_a$.

This ensures that $p_a \geq p_A$.

The h_i are all Bernoulli random variables. In a Bayesian setting, the retailer's knowledge of information can be described by the prior family which is a Beta distribution $B(r_i, n_i)$ a density function of $h_i^{r_i - 1}(1 - h_i)^{n_i - r_i - 1}$ with an expectation of(r_i / n_i). At any point, the retailers' belief about the acceptance probabilities of the offers $p_a > p_A$ is by the parameters (r_a, n_a, r_A, n_A).

r_a, n_a = the parameters of the Beta distribution describing ones belief of p_a. If offer a is accepted, the parameters get updated to $r_a + 1$, $n_a + 1$, while when it is rejected they get updated to r_a, $n_a + 1$. Thus one could reinterpret these as

- r_a = the number of people already accepted Offer a;
- n_a = the number of people already offered Offer a.

r_A, n_A = the parameters of the Beta distribution describing ones belief of p_A. If an offer A is accepted, the parameters get updated to $r_A + 1$, $n_A + 1$, while it is rejected but the customer would have accepted Offer a they get updated to r_A, $n_A + 1$ and the r_a, n_a remain unchanged. Thus one could reinterpret these as if

- r_A = the number of people already accepted Offer A, given they would accept offer a;
- n_A = the number of people who have been offered Offer A, who would have accepted offer a.

In both of the cases, observe that $n_i \geq r_i$ for $i = a, A$ must hold.

The selection of which product to offer is done by referring to information obtained based on the history of acceptance and rejections of the products. It is considered to be a "learning" model as it is basing its decision of which product to present to a prospective customer on past information.

In this chapter, we look at offering the two variants of a product to two different populations. Each population has its own probability of acceptance of Offer a and Offer A. Hence the gains from accepting the different offers will differ as well. We consider a number of variations of the problem, depending on whether the accepted probabilities of the offers are known or not, whether the type of customer is identifiable or not, and if not, how many questions one can ask. Table 5.1 shows which section each of these problems is considered in.

Section	Probability of accepting Offer a	Probability of accepting Offer A	Knowledge of customer type	Number of questions
5.1	Known	Known	Known	0
5.2	Known	Unknown	Known	0
5.3	Known	Known	Unknown	1
5.4	Known	Unknown	Unknown	1
5.5	Known	Known	Unknown	2
5.6	Known	Unknown	Unknown	2

Table 5.1: List of the problems considered in Chapter 5

5.1 Known probabilities of accepts for the two offers complete knowledge of origin of customer

In this section, we consider the population is heterogeneous, with 2 different populations which we refer to as Population 1 and Population 2. Each population has a different but known acceptance rate for the two offers available which are Offer a and Offer A. Both populations also have a different profit/gain from accepting offers a and A. All this is denoted as follows:

p_1 = Probability of accept for offer a by Population 1

p_2 = Probability of accept for offer a by Population 2

g_1 = Gain of acceptance of offer a by Population 1

g_2 = Gain of acceptance of offer a by Population 2

P_1 = Probability of accept for offer A by Population 1

P_2 = Probability of accept for offer A by Population 2

G_1 = Gain of acceptance of offer A by Population 1

G_2 = Gain of acceptance of offer A by Population 2

o_1 = Percentage of Population 1 in the entire population set

o_2 = Percentage of Population 2 in the entire population set

$o_1 + o_2 = 1$ where $i = 1, 2$ (population).

Note that $g_1 < G_1$ and $g_2 < G_2$ as the profit for Offer a is lower than that Offer A. Both a and A are variants of an offer that are ordered so their attractiveness to the customers hence the acceptance probability of that variant is monotonically decreasing while the variant's profitability to the lender is monotonically increasing.

We do not have any information on what happens if Offer A is rejected but it is our belief that Offer a would have been accepted. Here, the retailer is equipped with the information of which population a customer comes from. This problem splits into two independent problems since the probability of acceptance for both offers for population 1 and 2 are known. If

$v_i(r_i, n_i)$ = expected maximum total profit per person from population i applies then,

$$v_i(r_i, n_i) = \max\{o_i p_i g_i, o_i P_i G_i\}$$

Hence, the optimal course of action is the following:

(1) Choose Offer a for Population 1 if $p_1 g_1 > P_1 G_1$; and Offer A otherwise.

(2) Choose Offer a for Population 2 if $p_2 g_2 > P_2 G_2$; and Offer A otherwise.

The solution here is obvious since the retailer knows the type of customer he is faced with the different accept rates and gains from accepting any of the two offers for the two populations. Hence the optimal offer strategy is to select the offer that generates the largest expected future profit from that applicant.

5.2 Known Acceptance Probability of Offer a (p_i) but Unknown Acceptance Probability of Offer A (P_i), Complete Knowledge of the Origin of Customer

Consider then a situation where the acceptance probability for only one offer (the base offer, offer a) is known. Again, we assume the retailer has complete knowledge of which population the customer comes from. This splits into again two independent problems. Since we know which population the customer belongs to, then only information on that population, named Population i, affects the decision if a person from i applies. Hence that decision only affects next time when someone from i applies. Note that each of them is an easier version of the problem from Chapter 3 of the thesis. So, let

$v_i(r_i, n_i)$ = Expected maximum total future profit if person from population i has just applied

β_i = another chance there is another applicant from population i

The optimality equation is as follows:

$$v_i(r_i,n_i)=\max\left\{\begin{array}{c} p_i g_i+\beta_i\big(p_i v_i(r_i+1,n_i+1)+(1-p_i)v_i(r_i,n_i)\big), \\ \left(p_i\frac{r_i}{n_i}\right)G_i+\beta_i\left(p_i\frac{r_i}{n_i}v_i(r_i+1,n_i+1)+p_i\left(1-\frac{r_i}{n_i}\right)v_i(r_i,n_i+1)+(1-p_i)v_i(r_i,n_i)\right)\end{array}\right\}.$$

where $\beta_i=\dfrac{\beta o_i}{1-\beta(1-o_i)}$

Note that

$v(i, r_1, n_1, r_2, n_2) =$ Expected maximum total future profit of person from population i

$$v(i, r_1, n_1, r_2, n_2) = p_i g_i + \beta(o_1(p_1 v(1, r_1 + 1, n_1 + 1, r_2, n_2) + (1 - p_i)v(1, r_1, n_1 + 1, r_2, n_2)) +$$
$$+ \beta(o_2(p_1 v(2, r_1 + 1, n_1 + 1, r_2, n_2) + (1 - p_1)v(2, r_1, n_1 + 1, r_2, n_2))) + \ldots$$

For example, the chance that there is one applicant from Population 1 is:

$$\begin{aligned}
&\beta o_1 + \beta^2 o_2 o_1 + \beta^3 o_2^2 o_1 + \ldots \\
&= \beta o_1 (1 + \beta o_2 + \beta^2 o_2^2 + \ldots) \\
&= \frac{\beta o_1}{1 - \beta o_2}.
\end{aligned}$$

Hence the probability that there is another customer is from Population i is given by $\beta_i = \frac{\beta o_i}{1 - \beta(1 - o_i)}$. The first term in each equation is the probability that the next customer from population i, where $i = 1, 2$, will accept offer a or offer A multiplied for the gain of each population to the lender. The remaining terms of the equations are subject to the chance of β that there will be another customer. The first term corresponds to the customer from population i accepting offer a or offer A (depending on the equation) and the last term of each equation describes the offer being rejected detailing the different ways it could happen.

If Offer A is rejected, we do not have any information on whether they would accept Offer a, but it is our belief that Offer a would be accepted. We denote the chance of rejecting Offer A is $\left(1 - p_i \frac{r_i}{n_i}\right)$, $i = 1, 2$ where the probability of accepting Offer a is p_i. So,

$$\text{Prob}(\text{accept Offer } a | \text{reject Offer } A) = \frac{\text{P}(\text{accept Offer } a \text{ and reject Offer } A)}{\text{P}(\text{reject offer } A)}$$
$$= \frac{p_i\left(1 - \frac{r_i}{n_i}\right)}{1 - p_i \frac{r_i}{n_i}}$$

$$\therefore \text{Prob (accept Offer } a \text{ and reject Offer } A) = \frac{p_i\left(1-\frac{r_i}{n_i}\right)}{\left(1-p_i\frac{r_i}{n_i}\right)} \times \left(1-p_i\frac{r_i}{n_i}\right)$$

$$= p_i\left(1-\frac{r_i}{n_i}\right).$$

We once again consider a variation of the problem of offers a and A where the retailer incurs a cost of $\beta_i((1-p_i)v_i(r_i, n_i))$ whenever an offer is made to a customer in a state denoted by (r_1, n_1, r_2, n_2) regardless of which population the customer originates from. Note that this cost is not dependant on the offer made, and hence does not effect the optimal action. So, the optimal policy for $v_i(r_i, n_i)$ is the same for $\widetilde{V}_i(r_i, n_i)$ where

$$\widetilde{V}_i(r_i, n_i) = p_i\beta_i \max\begin{cases} B_1 + \widetilde{V}_i(r_i+1, n_i+1), \\ \frac{r_i}{n_i}B_2 + \frac{r_i}{n_i}\widetilde{V}_i(r_i+1, n_i+1) + \left(1-\frac{r_i}{n_i}\right)\widetilde{V}_i(r_i, n_i+1). \end{cases} \tag{7}$$

where $B_1 = \frac{g_i}{\beta_i}$ and $B_2 = \frac{G_i}{\beta_i}$

The first equation on the right calculates the expected total future profit from population i if offer a is accepted, while the second equation is for the expected total future profit from population i if offer A is accepted. The next step is to prove the monotonic characteristics of the optimal equations stated in **(7)**.

Lemma 5.1:

$\widetilde{V}_i(r_i, n_i)$ is (i) non decreasing in r_i.

(ii) non increasing in n_i.

Proof

The proof is similar to the one presented in Lemma 3.1. For details, please refer to **section 3.2,** Chapter 3.

5.3 Known Probabilities of Acceptance for Both Offers, Incomplete Knowledge of the Origin of Customer – One Question

Under this condition, a question is allowed to gain some information to which population does the current applicant belong and hence help decide of which offer to extend. There are two responses to the questions chosen to be asked; "yes" or "no".

We introduce the following new variables:

q_1 = Probability of Population 1 saying "Yes" to the question

q_2 = Probability of Population 2 saying "Yes" to the question

When the question is asked, whether Offer a or Offer A is extended will depend on the response to the question. So, one can phrase the following condition as follows:

$$\text{Prob}\left(\text{Population} = 1 \middle| \text{Answer} = \text{"yes"}\right) = \frac{P(y|1)P(1)}{P(y)} = \frac{q_1 o_1}{q_1 o_1 + q_2 o_2} \text{ and}$$

$$\text{Prob}\left(\text{Population} = 2 \middle| \text{Answer} = \text{"yes"}\right) = \frac{P(y|2)P(2)}{P(y)} = \frac{q_2 o_2}{q_1 o_1 + q_2 o_2}.$$

There is no learning in this problem as the retailer has complete knowledge about the acceptance rate for all the offers from each population. Thus one makes the same decisions for all applicants irrespective of how many more are to be considered.

$v(y)$ = Expected maximum total profit from this applicant if the response to the question is "yes",

$v(n)$ = Expected maximum total profit from this applicant if the response to the question is "no".

If answer given is "yes" (y):

$$v(y)=\max\begin{cases}\left(\dfrac{q_1o_1}{q_1o_1+q_2o_2}\right)p_1g_1+\left(\dfrac{q_2o_2}{q_1o_1+q_2o_2}\right)p_2g_2,\\ \left(\dfrac{q_1o_1}{q_1o_1+q_2o_2}\right)P_1G_1+\left(\dfrac{q_2o_2}{q_1o_1+q_2o_2}\right)P_2G_2.\end{cases}$$

Hence the condition where Offer a can be chosen can be expressed in the form of an inequality of:

$$q_1o_1p_1g_1+q_2o_2p_2g_2>q_1o_1P_1G_1+q_2o_2P_2G_2$$

$$\Leftrightarrow q_1o_1(p_1g_1-P_1G_1)>q_2o_2(P_2G_2-p_2g_2)$$

$$\Leftrightarrow\left(\frac{q_1}{q_2}\right)\left(\frac{o_1}{1-o_1}\right)>\frac{P_2G_2-p_2g_2}{p_1g_1-P_1G_1}.$$

If answer given is "no" (n):

$$v(n)=\max\begin{cases}\left(\dfrac{(1-q_1)o_1}{(1-q_1)o_1+(1-q_2)o_2}\right)p_1g_1+\left(\dfrac{(1-q_2)o_2}{(1-q_1)o_1+(1-q_2)o_2}\right)p_2g_2,\\ \left(\dfrac{(1-q_1)o_1}{(1-q_1)o_1+(1-q_2)o_2}\right)P_1G_1+\left(\dfrac{(1-q_2)o_2}{(1-q_1)o_1+(1-q_2)o_2}\right)P_2G_2.\end{cases}$$

The condition where Offer a can be chosen can be expressed in the form of an inequality of:

$$(1-q_1)o_1p_1g_1+(1-q_2)o_2p_2g_2>(1-q_1)o_1P_1G_1+(1-q_2)o_2P_2G_2$$

$$\Leftrightarrow(1-q_1)o_1(p_1g_1-P_1G_1)>(1-q_2)o_2(P_2G_2-p_2g_2)$$

$$\Leftrightarrow\left(\frac{1-q_1}{1-q_2}\right)\left(\frac{o_1}{o_2}\right)>\frac{P_2G_2-p_2g_2}{p_1g_1-P_1G_1}$$

$$\Leftrightarrow\left(\frac{1-q_1}{1-q_2}\right)\left(\frac{o_1}{1-o_1}\right)>\frac{P_2G_2-p_2g_2}{p_1g_1-P_1G_1}.$$

Hence the decision on the optimal offer strategy will depend on the order of the following 3 ratios: $\frac{q_1}{q_2}$, $\frac{1-q_1}{1-q_2}$ and $\frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$.

Let

$v(r_1, n_1, r_2, n_2)$ = profit from an applicant before they answer the question

So

$$v(r_1, n_1, r_2, n_2) = (q_1o_1 + q_2o_2)v(y) + ((1-q_1)o_1 + (1-q_2)o_2)v(n)$$

Assuming that both

$$\left(\frac{q_1}{q_2}\right) > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right) \text{ and}$$

$$\left(\frac{1-q_1}{1-q_2}\right) > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right),$$

we look at the possible conditions and the corresponding optimal offer strategies:

(i) If we define the "yes" answer to be positive response to a question (as defined by the user) which is the case where $\frac{q_1}{q_2} > \frac{1-q_1}{1-q_2}$, then the possible strategy reduces to :

(a) extend Offer a to all, or

(b) extend Offer A to all, or

(c) extend Offer a to those applicants who answer "yes" and Offer A to those who answer "no".

(ii) If we define the "no" answer to be the negative response to a question (as defined by the user) which is the case where $\frac{1-q_1}{1-q_2} > \frac{q_1}{q_2}$, then the possible strategy reduces to:

(a) extend Offer a to all, or

(b) extend Offer A to all, or

(c) extend Offer a to those applicants who answer "no" and Offer A to those who answer "yes".

5.4 Known Probability of Acceptance for Offer *a* (p_i) but Unknown Probabilities of Acceptance for Offer *A* (P_i), Incomplete Knowledge of Origin of Customer – One Question

We introduce

$h_1, h_2 =$ Probability of Population 1(2) $\{\text{accept Offer } A \,|\, \text{accept Offer } a\}$

h_i are all Bernoulli random variables. In a Bayesian setting, we describe the retailer's knowledge of information using a Beta distribution $B(r_i, n_i)$ where the density function is $h_i^{r_i-1}(1-h_i)^{n_i-r_i-1}$ with an expectation of (r_i / n_i). You can interpret this as:

$r_1 =$ the number of applicants from Population 1 who have accepted the offer

$n_1 =$ the number of applicants from Population 1 who have been given an offer

$r_2 =$ the number of applicants from Population 2 who have accepted the offer

$n_2 =$ the number of applicants from Population 2 who have been given an offer

$v(y, r_1, n_1, r_2, n_2) =$ Maximum expected total future profit if the applicant replies "yes" to question

$v(n, r_1, n_1, r_2, n_2) =$ Maximum expected total future profit if the applicant replies "no" to question

We define

$v(r_1, n_1, r_2, n_2) =$ Profit from an applicant before answering the question

and this is calculated as follows:

$$v(r_1, n_1, r_2, n_2) = (q_1 o_1 + q_2 o_2) v(y, r_1, n_1, r_2, n_2) + ((1-q_1)o_1 + (1-q_2)o_2) v(n, r_1, n_1, r_2, n_2).$$

The optimality equations are as follows:

If response is "yes" (y):

$$v(y,r_1,n_1,r_2,n_2)=\max\begin{cases}\frac{q_1o_1}{q_1o_1+q_2o_2}p_1g_1+\frac{q_2o_2}{q_1o_1+q_2o_2}p_2g_2+\beta v(r_1,n_1,r_2,n_2),\\ \frac{q_1o_1}{q_1o_1+q_2o_2}p_1\frac{r_1}{n_1}G_1+\frac{q_2o_2}{q_1o_1+q_2o_2}p_2\frac{r_2}{n_2}G_2+\\ \beta(\frac{q_1o_1}{q_1o_1+q_2o_2}\{p_1\frac{r_1}{n_1}v(r_1+1,n_1+1,r_2,n_2)+p_1\left(1-\frac{r_1}{n_1}\right)v(r_1,n_1+1,r_2,n_2)+\\ (1-p_1)v(r_1,n_1,r_2,n_2)\}+\frac{q_2o_2}{q_1o_1+q_2o_2}\{p_2\frac{r_2}{n_2}v(r_1,n_1,r_2+1,n_2+1)+\\ p_2\left(1-\frac{r_2}{n_2}\right)v(r_1,n_1,r_2,n_2+1)+(1-p_2)v(r_1,n_1,r_2,n_2)\}).\end{cases}$$

If the response is a "no" (n):

$$v(n,r_1,n_1,r_2,n_2)=\max\begin{cases}\frac{(1-q_1)o_1}{(1-q_1)o_1+(1-q_2)o_2}p_1g_1+\frac{(1-q_2)o_2}{(1-q_1)o_1+(1-q_2)o_2}p_2g_2+\beta v(r_1,n_1,r_2,n_2),\\ \frac{(1-q_1)o_1}{(1-q_1)o_1+(1-q_2)o_2}p_1\frac{r_1}{n_1}G_1+\frac{(1-q_2)o_2}{(1-q_1)o_1+(1-q_2)o_2}p_2\frac{r_2}{n_2}G_2+\\ \beta(\frac{(1-q_1)o_1}{(1-q_1)o_1+(1-q_2)o_2}\{p_1\frac{r_1}{n_1}v(r_1+1,n_1+1,r_2,n_2)+p_1\left(1-\frac{r_1}{n_1}\right)v(r_1,n_1+1,r_2,n_2)+\\ (1-p_1)v(r_1,n_1,r_2,n_2)\}+\frac{(1-q_2)o_2}{(1-q_1)o_1+(1-q_2)o_2}\{p_2\frac{r_2}{n_2}v(r_1,n_1,r_2+1,n_2+1)+\\ p_2\left(1-\frac{r_2}{n_2}\right)v(r_1,n_1,r_2,n_2+1)+(1-p_2)v(r_1,n_1,r_2,n_2)\}).\end{cases}$$

Note that the equations for the $v(y,r_1,n_1,r_2,n_2)$ and $v(n,r_1,n_1,r_2,n_2)$ are almost identical except for the term $(1-q_1)$ in $v(n,r_1,n_1,r_2,n_2)$that replaces the q_1 in $v(n,r_1,n_1,r_2,n_2)$ with $(1-q_1)>0$. Hence if the monotonicity is proven for $v(y,r_1,n_1,r_2,n_2)$, the monotonicity will hold also for $v(n,r_1,n_1,r_2,n_2)$.

Using value iterations as we know that $v^m(r_1,n_1,r_2,n_2)$ converges to $v(r_1,n_1,r_2,n_2)$ where $v^m(r_1,n_1,r_2,n_2)$ is defined by:

if response is "yes" (y):

$$v^{m+1}(y, r_1, n_1, r_2, n_2) = \frac{\beta}{q_1 o_1 + q_2 o_2} \max \begin{cases} A_1 + A_2 + (q_1 o_1 + q_2 o_2) v^m(r_1, n_1, r_2, n_2), \\ \frac{r_1}{n_1} B_1 + \frac{r_1}{n_1} q_1 o_1 p_1 v^m(r_1 + 1, n_1 + 1, r_2, n_2) + \\ \left(1 - \frac{r_1}{n_1}\right) q_1 o_1 p_1 v^m(r_1, n_1 + 1, r_2, n_2) + \frac{r_2}{n_2} B_2 + \\ \frac{r_2}{n_2} q_2 o_2 p_2 v^m(r_1, n_1, r_2 + 1, n_2 + 1) + \\ \left(1 - \frac{r_2}{n_2}\right) q_2 o_2 p_2 v^m(r_1, n_1, r_2, n_2 + 1) + \\ q_1 o_1 (1 - p_1) v^m(r_1, n_1, r_2, n_2) + q_2 o_2 (1 - p_2) v^m(r_1, n_1, r_2, n_2). \end{cases}$$

where $A_1 = \frac{q_1 o_1 p_1 g_1}{\beta}$, $A_2 = \frac{q_2 o_2 p_2 g_2}{\beta}$, $B_1 = \frac{q_1 o_1 p_1 G_1}{\beta}$ and $B_2 = \frac{q_2 o_2 p_2 G_2}{\beta}$.

If the response is a "no" (n):

$$v^{m+1}(n, r_1, n_1, r_2, n_2) = \frac{\beta}{(1-q_1)o_1 + (1-q_2)o_2} \max \begin{cases} C_1 + C_2 + ((1-q_1)\, o_1 + (1-q_2)\, o_2) v^m(r_1, n_1, r_2, n_2), \\ \frac{r_1}{n_1} D_1 + \frac{r_1}{n_1}(1-q_1)\, o_1 p_1 v^m(r_1 + 1, n_1 + 1, r_2, n_2) + \\ \left(1 - \frac{r_1}{n_1}\right)(1-q_1)\, o_1 p_1 v^m(r_1, n_1 + 1, r_2, n_2) + \frac{r_2}{n_2} D_2 + \\ \frac{r_2}{n_2}(1-q_2)\, o_2 p_2 v^m(r_1, n_1, r_2 + 1, n_2 + 1) + \\ \left(1 - \frac{r_2}{n_2}\right)(1-q_2)\, o_2 p_2 v'^m(r_1, n_1, r_2, n_2 + 1) + \\ (1-q_1)(1-p_1)\, v^m(r_1, n_1, r_2, n_2) + (1-q_2)(1-p_2)\, v^m(r_1, n_1, r_2, n_2)\}). \end{cases}$$

where $C_1 = \frac{(1-q_1)o_1 p_1 g_1}{\beta}$, $C_2 = \frac{(1-q_2)o_2 p_2 g_2}{\beta}$, $D_1 = \frac{(1-q_1)o_1 p_1 G_1}{\beta}$ and $D_2 = \frac{(1-q_2)o_2 p_2 G_2}{\beta}$.

with $v^0(r_1, n_1, r_2, n_2) = 0$ which means trivially $v^1(r_1, n_1, r_2, n_2) \geq v^0(r_1, n_1, r_2, n_2)$. We use this induction to prove $v^m(r_1, n_1, r_2, n_2) \geq v^{m-1}(r_1, n_1, r_2, n_2)$.

Since $\max\{a,b\}-\max\{c,d\}\geq\min\{a-c,b-d\}$, so

$$v^{m+1}(r_1,n_1,r_2,n_2)-v^m(r_1,n_1,r_2,n_2)\geq\beta\{v^m(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2)\}>0$$

So $v^{m+1}(r_1,n_1,r_2,n_2)$ is a monotone increasing sequence that is bounded by $\frac{\beta}{1-\beta}\max\{g_1,g_2,G_1,G_2\}$ as $v^{m+1}(r_1,n_1,r_2,n_2)\leq\frac{\beta}{1-\beta}\max\{g_1,g_2,G_1,G_2\}$. So the iterates of the value iteration will converge to a $v(r_1,n_1,r_2,n_2)$.

We now use the iterates of the value iteration to prove

$v^{m+1}(y,r_1,n_1,r_2,n_2)-v^m(y,r_1,n_1,r_2,n_2)>0$ and

$v^{m+1}(n,r_1,n_1,r_2,n_2)-v^m(n,r_1,n_1,r_2,n_2)>0$.

$$\begin{aligned}
&v^{m+1}(y,r_1,n_1,r_2,n_2)-v^m(y,r_1,n_1,r_2,n_2)\\
&=\max\begin{cases}
(q_1o_1+q_2o_2)\left(v^m(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2)\right),\\
\frac{r_1}{n_1}q_1o_1p_1\left(v^m(r_1+1,n_1+1,r_2,n_2)-v^{m-1}(r_1+1,n_1+1,r_2,n_2)\right)+\\
\left(1-\frac{r_1}{n_1}\right)q_1o_1p_1\left(v^m(r_1,n_1+1,r_2,v_2)-v^{m-1}(r_1,n_1+1,r_2,n_2)\right)+\\
q_1o_1(1-p_1)\left(v^m(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2)\right)+\\
\frac{r_2}{n_2}q_2o_2p_2\left(v^m(r_1,n_1,r_2+1,n_2+1)-v^{m-1}(r_1,n_1,r_2+1,n_2+1)\right)+\\
\left(1-\frac{r_2}{n_2}\right)\left(v^m(r_1,n_1,r_2,n_2+1)-v^{m-1}(r_1,n_1,r_2,n_2+1)\right)+\\
q_2o_2(1-p_2)\left(v^m(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2)\right).
\end{cases}\\
&>0.
\end{aligned}$$

If the response is "no", hence

$$v^{m+1}(n, r_1, n_1, r_2, n_2) - v^m(n, r_1, n_1, r_2, n_2)$$

$$= \max\begin{cases} ((1-q_1)o_1 + (1-q_2)o_2)\left(v^m(r_1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right), \\ \frac{r_1}{n_1}(1-q_1)o_1 p_1\left(v^m(r_1+1, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right) + \\ \left(1-\frac{r_1}{n_1}\right)(1-q_1)o_1 p_1\left(v^m(r_1, n_1+1, r_2, v_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ (1-q_1)o_1(1-p_1)\left(v^m(r_1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) + \\ \frac{r_2}{n_2}(1-q_2)o_2 p_2\left(v^m(r_1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right) + \\ \left(1-\frac{r_2}{n_2}\right)(1-q_2)o_2 p_2\left(v^m(r_1, n_1, r_2, n_2+1) - v^{m-1}(r_1, n_1, r_2, n_2+1)\right) + \\ (1-q_2)o_2(1-p_2)\left(v^m(r_1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right). \end{cases}$$

$$> 0.$$

If this condition holds for m, it will hold for $m+1$. We use this result to prove some general results in Lemma 5.2.

Lemma 5.2:

$v^{m+1}(r_1, n_1, r_2, n_2)$ is (i) non decreasing in r_i.

(ii) non increasing in n_i.

Proof

Recall the optimality equations can be rewritten as:

If response is "yes" (y):

$$v^m(y, r_1, n_1, r_2, n_2) = \frac{\beta}{q_1 o_1 + q_2 o_2}\max\begin{cases} A_1 + A_2 + (q_1 o_1 + q_2 o_2)v^{m-1}(r_1, n_1, r_2, n_2), \\ \frac{r_1}{n_1}B_1 + \frac{r_1}{n_1}q_1 o_1 p_1 v^{m-1}(r_1+1, n_1+1, r_2, n_2) + \\ \left(1-\frac{r_1}{n_1}\right)q_1 o_1 p_1 v^{m-1}(r_1, n_1+1, r_2, n_2) + \frac{r_2}{n_2}B_2 + \\ \frac{r_2}{n_2}q_2 o_2 p_2 v^{m-1}(r_1, n_1, r_2+1, n_2+1) + \\ \left(1-\frac{r_2}{n_2}\right)q_2 o_2 p_2 v^{m-1}(r_1, n_1, r_2, n_2+1) + \\ q_1 o_1(1-p_1)v^{m-1}(r_1, n_1, r_2, n_2) + q_2 o_2(1-p_2)v^{m-1}(r_1, n_1, r_2, n_2). \end{cases}$$

where $A_1 = \dfrac{q_1 o_1 p_1 g_1}{\beta}$, $A_2 = \dfrac{q_2 o_2 p_2 g_2}{\beta}$, $B_1 = \dfrac{q_1 o_1 p_1 G_1}{\beta}$ and $B_2 = \dfrac{q_2 o_2 p_2 G_2}{\beta}$.

If the response is a "no" (n):

$$v^m(n, r_1, n_1, r_2, n_2) = \frac{\beta}{(1-q_1)o_1 + (1-q_2)o_2} \max \begin{cases} C_1 + C_2 + ((1-q_1)\, o_1 + (1-q_2)\, o_2) v^{m-1}(r_1, n_1, r_2, n_2), \\ \frac{r_1}{n_1} D_1 + \frac{r_1}{n_1}(1-q_1)\, o_1 p_1 v^{m-1}(r_1+1, n_1+1, r_2, n_2) + \\ \left(1 - \frac{r_1}{n_1}\right)(1-q_1)\, o_1 p_1 v^{m-1}(r_1, n_1+1, r_2, n_2) + \frac{r_2}{n_2} D_2 + \\ \frac{r_2}{n_2}(1-q_2)\, o_2 p_2 v^{m-1}(r_1, n_1, r_2+1, n_2+1) + \\ \left(1 - \frac{r_2}{n_2}\right)(1-q_2)\, o_2 p_2 v^{m-1}(r_1, n_1, r_2, n_2+1) + \\ (1-q_1)(1-p_1)\, v^{m-1}(r_1, n_1, r_2, n_2) + (1-q_2)(1-p_2)\, v^{m-1}(r_1, n_1, r_2, n_2)\}). \end{cases}$$

where $C_1 = \dfrac{(1-q_1)o_1 p_1 g_1}{\beta}$, $C_2 = \dfrac{(1-q_2)o_2 p_2 g_2}{\beta}$, $D_1 = \dfrac{(1-q_1)o_1 p_1 G_1}{\beta}$ and $D_2 = \dfrac{(1-q_2)o_2 p_2 G_2}{\beta}$

The definition for $v^{m+1}(r_1, n_1, r_2, n_2)$ is maintained as profit from an applicant before answering the question and is calculated as follows:

$$v^{m+1}(r_1, n_1, r_2, n_2) = (q_1 o_1 + q_2 o_2) v^m(y, r_1, n_1, r_2, n_2) + ((1-q_1)o_1 + (1-q_2)o_2) v^m(n, r_1, n_1, r_2, n_2).$$

Next, we proceed with proving result (i):

(i) $v^m(r_1+1, n_1, r_2, n_2) - v^m(r_1, n_1, r_2, n_2) \geq 0$,

$v^m(r_1, n_1, r_2+1, n_2) - v^m(r_1, n_1, r_2, n_2) \geq 0$.

Assume this is true for $m-1$ with $v^0(r_1, n_1, r_2, n_2) = 0$.

We take the optimality equation from the "yes" category. So,

$$v^m(y, r_1+1, n_1, r_2, n_2) - v^m(y, r_1, n_1, r_2, n_2)$$

$$\geq \min\begin{cases} (q_1 o_1 + q_2 o_2)\left(v^{m-1}(r_1+1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right), \\ \left(\frac{1}{n_1}\right)\left(B_1 + q_1 o_1 p_1\left(v^{m-1}(r_1+2, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right)\right) + \\ \left(\frac{r_1}{n_1}\right) q_1 o_1 p_1\left(v^{m-1}(r_1+2, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right) + \\ \left(1 - \frac{r_1}{n_1}\right) q_1 o_1 p_1\left(v^{m-1}(r_1+1, n_1+1, r_2, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ \left(\frac{r_2}{n_2}\right) q_2 o_2 p_2\left(v^{m-1}(r_1+1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right) + \\ \left(1 - \frac{r_2}{n_2}\right) q_2 o_2 p_2\left(v^{m-1}(r_1+1, n_1, r_2, n_2+1) - v^{m-1}(r_1, n_1, r_2, n_2+1)\right) + \\ q_1 o_1 (1-p_1)\left(v^{m-1}(r_1+1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) + \\ q_2 o_2 (1-p_2)\left(v^{m-1}(r_1+1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right). \end{cases}$$

$$> 0\,.$$

Trivially the same proof will show that $v^m(n, r_1+1, n_1, r_2, n_2) \geq v^m(n, r_1, n_1, r_2, n_2)$.

$$v^m(n, r_1+1, n_1, r_2, n_2) - v^m(n, r_1, n_1, r_2, n_2)$$

$$\geq \min\begin{cases} ((1-q_1)o_1 + (1-q_2)o_2)\left(v^{m-1}(r_1+1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right), \\ \left(\frac{1}{n_1}\right)\left(D_1 + (1-q_1)o_1 p_1\left(v^{m-1}(r_1+2, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right)\right) + \\ \left(\frac{r_1}{n_1}\right)(1-q_1)o_1 p_1\left(v^{m-1}(r_1+2, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right) + \\ \left(1 - \frac{r_1}{n_1}\right)(1-q_1)o_1 p_1\left(v^{m-1}(r_1+1, n_1+1, r_2, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ \left(\frac{r_2}{n_2}\right)(1-q_2)o_2 p_2\left(v^{m-1}(r_1+1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right) + \\ \left(1 - \frac{r_2}{n_2}\right)(1-q_2)o_2 p_2\left(v^{m-1}(r_1+1, n_1, r_2, n_2+1) - v^{m-1}(r_1, n_1, r_2, n_2+1)\right) + \\ (1-q_1)o_1(1-p_1)\left(v^{m-1}(r_1+1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) + \\ (1-q_2)o_2(1-p_2)\left(v^{m-1}(r_1+1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right). \end{cases}$$

$$> 0\,.$$

Hence $v^m(r_1+1, n_1, r_2, n_2) \geq v^m(r_1, n_1, r_2, n_2)$.

Also $v^m(y, r_1, n_1, r_2+1, n_2) - v^m(y, r_1, n_1, r_2, n_2)$

$$\ge \min\begin{cases} (q_1 o_1 + q_2 o_2)\left(v^{m-1}(r_1, n_1, r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right), \\ \left(\frac{r_1}{n_1}\right) q_1 o_1 p_1 \left(v^{m-1}(r_1+1, n_1+1, r_2+1, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right) + \\ \left(1-\frac{r_1}{n_1}\right) q_1 o_1 p_1 \left(v^{m-1}(r_1, n_1+1, r_2+1, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ \left(\frac{1}{n_2}\right)\left(B_2 + q_2 o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2+2, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right)\right) \\ \left(\frac{r_2}{n_2}\right) q_2 o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2, +2\, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right) + \\ \left(1-\frac{r_2}{n_2}\right) q_2 o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1, r_2, n_2+1)\right) + \\ q_1 o_1 (1-p_1)\left(v^{m-1}(r_1, n_1\; r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) + \\ q_2 o_2 (1-p_2)\left(v^{m-1}(r_1, n_1\; r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) \end{cases}$$

> 0.

Trivially the same proof will show that $v^m(n, r_1, n_1, r_2+1, n_2) \ge v^m(n, r_1, n_1, r_2, n_2)$.

$v^m(n, r_1, n_1, r_2+1, n_2) - v^m(n, r_1, n_1, r_2, n_2)$

$$\ge \min\begin{cases} ((1-q_1)o_1 + (1-q_2)o_2)\left(v^{m-1}(r_1, n_1, r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right), \\ \left(\frac{r_1}{n_1}\right)(1-q_1) o_1 p_1 \left(v^{m-1}(r_1+1, n_1+1, r_2+1, n_2) - v^{m-1}(r_1+1, n_1+1, r_2, n_2)\right) + \\ \left(1-\frac{r_1}{n_1}\right)(1-q_1) o_1 p_1 \left(v^{m-1}(r_1, n_1+1, r_2+1, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ \left(\frac{1}{n_2}\right)\left(D_2 + (1-q_2) o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2+2, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right)\right) \\ \left(\frac{r_2}{n_2}\right)(1-q_2) o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2, +2\, n_2+1) - v^{m-1}(r_1, n_1, r_2+1, n_2+1)\right) + \\ \left(1-\frac{r_2}{n_2}\right)(1-q_2) o_2 p_2 \left(v^{m-1}(r_1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1, r_2, n_2+1)\right) + \\ (1-q_1) o_1 (1-p_1)\left(v^{m-1}(r_1, n_1\; r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) + \\ (1-q_2) o_2 (1-p_2)\left(v^{m-1}(r_1, n_1\; r_2+1, n_2) - v^{m-1}(r_1, n_1, r_2, n_2)\right) \end{cases}$$

Hence $v^m(r_1, n_1, r_2+1, n_2) \ge v^m(r_1, n_1, r_2, n_2)$. Since result (i) holds for m, so it will hold for $m+1$. By induction, this holds then for all m.

We use the same technique to prove result (ii):

(ii) $$v^m(r_1, n_1, r_2, n_2) - v^m(r_1, n_1+1, r_2, n_2) \geq 0,$$

$$v^m(r_1, n_1, r_2, n_2) - v^m(r_1, n_1, r_2, n_2+1) \geq 0.$$

Assume that this is true for $m-1$ with $v^0(r_1, n_1, r_2, n_2) = 0$.

We take the equations from the "yes" response:

$$v^m(y, r_1, n_1, r_2, n_2) - v^m(y, r_1, n_1+1, r_2, n_2)$$

$$\geq \min \begin{cases} (q_1 o_1 + q_2 o_2)\left(v^{m-1}(r_1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right), \\ \left(\dfrac{r_1}{n_1(n_1+1)}\right)\left(B_1 + q_1 o_1 p_1\left(v^{m-1}(r_1+1, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+2, r_2, n_2)\right)\right) + \\ \left(\dfrac{r_1}{n_1+1}\right) q_1 o_1 p_1\left(v^{m-1}(r_1+1, n_1+1, r_2, n_2) - v^{m-1}(r_1+1, n_1+2, r_2, n_2)\right) + \\ \left(1 - \dfrac{r_1}{n_1}\right) q_1 o_1 p_1\left(v^{m-1}(r_1, n_1+1, r_2, n_2) - v^{m-1}(r_1, n_1+2, r_2, n_2)\right) + \\ \left(\dfrac{r_2}{n_2}\right) q_2 o_2 p_2\left(v^{m-1}(r_1, n_1, r_2+1, n_2+1) - v^{m-1}(r_1, n_1+1, r_2+1, n_2+1)\right) + \\ \left(1 - \dfrac{r_2}{n_2}\right) q_2 o_2 p_2\left(v^{m-1}(r_1, n_1, r_2, n_2+1) - v^{m-1}(r_1, n_1+1, r_2, n_2+1)\right) + \\ q_1 o_1 (1-p_1)\left(v^{m-1}(r_1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right) + \\ q_2 o_2 (1-p_2)\left(v^{m-1}(r_1, n_1\ r_2, n_2) - v^{m-1}(r_1, n_1+1, r_2, n_2)\right). \end{cases}$$

>0.

Trivially, this proof will show that $v^m(n, r_1, n_1, r_2, n_2) \geq v^m(n, r_1, n_1+1, r_2, n_2)$.

$$v^m(n,r_1,n_1,r_2,n_2)-v^m(n,r_1,n_1+1,r_2,n_2)$$

$$\geq \min\begin{cases} \left((1-q_1)o_1+(1-q_2)o_2\right)\left(v^{m-1}(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1+1,r_2,n_2)\right), \\ \left(\frac{r_1}{n_1(n_1+1)}\right)\left(D_1+(1-q_1)o_1p_1\left(v^{m-1}(r_1+1,n_1+1,r_2,n_2)-v^{m-1}(r_1+1,n_1+2,r_2,n_2)\right)\right)+ \\ \left(\frac{r_1}{n_1+1}\right)(1-q_1)o_1p_1\left(v^{m-1}(r_1+1,n_1+1,r_2,n_2)-v^{m-1}(r_1+1,n_1+2,r_2,n_2)\right)+ \\ \left(1-\frac{r_1}{n_1}\right)(1-q_1)o_1p_1\left(v^{m-1}(r_1,n_1+1,r_2,n_2)-v^{m-1}(r_1,n_1+2,r_2,n_2)\right)+ \\ \left(\frac{r_2}{n_2}\right)(1-q_2)o_2p_2\left(v^{m-1}(r_1,n_1,r_2+1,n_2+1)-v^{m-1}(r_1,n_1+1,r_2+1,n_2+1)\right)+ \\ \left(1-\frac{r_2}{n_2}\right)(1-q_2)o_2p_2\left(v^{m-1}(r_1,n_1,r_2,n_2+1)-v^{m-1}(r_1,n_1+1,r_2,n_2+1)\right)+ \\ (1-q_1)o_1(1-p_1)\left(v^{m-1}(r_1,n_1\,r_2,n_2)-v^{m-1}(r_1,n_1+1,r_2,n_2)\right)+ \\ (1-q_2)o_2(1-p_2)\left(v^{m-1}(r_1,n_1\,r_2,n_2)-v^{m-1}(r_1,n_1+1,r_2,n_2)\right). \end{cases}$$

>0.

Hence $v^m(r_1,n_1,r_2,n_2)\geq v^m(r_1,n_1+1,r_2,n_2)$.

$$v^m(y,r_1,n_1,r_2,n_2)-v^m(y,r_1,n_1,r_2,n_2+1)$$

$$\geq \min\begin{cases} (q_1o_1+q_2o_2)\left(v^{m-1}(r_1,n_1,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2+1)\right), \\ \left(\frac{r_1}{n_1}\right)q_1o_1p_1\left(v^{m-1}(r_1+1,n_1+1,r_2,n_2)-v^{m-1}(r_1+1,n_1+1,r_2,n_2+1)\right)+ \\ \left(1-\frac{r_1}{n_1}\right)q_1o_1p_1\left(v^{m-1}(r_1,n_1+1,r_2,n_2)-v^{m-1}(r_1,n_1+1,r_2,n_2+1)\right)+ \\ \left(\frac{r_2}{n_2(n_2+1)}\right)\left(B_2+q_2o_2p_2\left(v^{m-1}(r_1,n_1,r_2+1,n_2+1)-v^{m-1}(r_1,n_1,r_2+1,n_2+2)\right)\right) \\ \left(\frac{r_2}{n_2+1}\right)q_2o_2p_2\left(v^{m-1}(r_1,n_1,r_2+1,n_2+1)-v^{m-1}(r_1,n_1,r_2+1,n_2+2)\right)+ \\ \left(1-\frac{r_2}{n_2}\right)q_2o_2p_2\left(v^{m-1}(r_1,n_1,r_2,n_2+1)-v^{m-1}(r_1,n_1,r_2,n_2+2)\right)+ \\ q_1o_1(1-p_1)\left(v^{m-1}(r_1,n_1\,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2+1)\right)+ \\ q_2o_2(1-p_2)\left(v^{m-1}(r_1,n_1\,r_2,n_2)-v^{m-1}(r_1,n_1,r_2,n_2+1)\right) \end{cases}$$

>0.

The same proof will also show that $v^m(n, r_1, n_1, r_2, n_2) \ge v^m(n, r_1, n_1, r_2, n_2 + 1)$

$$v^m(n, r_1, n_1, r_2, n_2) \ge v^m(n, r_1, n_1, r_2, n_2 + 1)$$

$$\ge \min \begin{cases} ((1-q_1)o_1 + (1-q_2)o_2)(v^{m-1}(r_1, n_1, r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2 + 1)), \\ \left(\frac{r_1}{n_1}\right)(1-q_1)o_1 p_1 (v^{m-1}(r_1 + 1, n_1 + 1, r_2, n_2) - v^{m-1}(r_1 + 1, n_1 + 1, r_2, n_2 + 1)) + \\ \left(1 - \frac{r_1}{n_1}\right)(1-q_1)o_1 p_1 (v^{m-1}(r_1, n_1 + 1, r_2, n_2) - v^{m-1}(r_1, n_1 + 1, r_2, n_2 + 1)) + \\ \left(\frac{r_2}{n_2(n_2+1)}\right)(D_2 + (1-q_2)o_2 p_2 (v^{m-1}(r_1, n_1, r_2 + 1, n_2 + 1) - v^{m-1}(r_1, n_1, r_2 + 1, n_2 + 2))) \\ \left(\frac{r_2}{n_2 + 1}\right)(1-q_2)o_2 p_2 (v^{m-1}(r_1, n_1, r_2 + 1, n_2 + 1) - v^{m-1}(r_1, n_1, r_2 + 1, n_2 + 2)) + \\ \left(1 - \frac{r_2}{n_2}\right)(1-q_2)o_2 p_2 (v^{m-1}(r_1, n_1, r_2, n_2 + 1) - v^{m-1}(r_1, n_1, r_2, n_2 + 2)) + \\ (1-q_1)o_1(1-p_1)(v^{m-1}(r_1, n_1 r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2 + 1)) + \\ (1-q_2)o_2(1-p_2)(v^{m-1}(r_1, n_1 r_2, n_2) - v^{m-1}(r_1, n_1, r_2, n_2 + 1)) \end{cases}$$

>0.

Hence $v^m(r_1, n_1, r_2, n_2) \ge v^m(r_1, n_1, r_2, n_2 + 1)$. Since result (i) holds for m, so it will hold for $m+1$. By induction, this holds then for all m.

Therefore, we have proven the monotonicity of the optimality equations for this case as well. The optimality equations used here are similar to the ones used in Chapter 3 but we have redefined r_i and n_i to reflect the population from which the accepts and recipients of offers originate.

Theorem 5.2:

$v(r_1, n_1, r_2, n_2)$ is (i) non decreasing in r_i.

(ii) non increasing in n_i.

Proof

Let $m \to \infty$ as in Lemma 5.2.

5.4.1 Optimality Equations

In **Theorem 3.1**, it states that in the limit when $s = \infty$ where s is the number of offers accepted, the optimal policy is in the form of a function $r^*(n_2)$ where

(a) in (r_2, n_2) with $r_2 \leq r_2^*(n_2)$, variant 1 is chosen,

(b) in (r_2, n_2) with $r_2 > r_2^*(n_2)$, variant 2 is chosen.

The experimental results we will display in the next section suggest that for this case of unknown probability of accepts for offer for Population 2 and incomplete information of the origin of the applicant, there exist an optimal policy in the form of a function of r as follows:

For Population 1:

(a) in (r_1, n_1, r_2, n_2) with a response of "Yes" and $r_1 \leq r_1^*(y, n_1, r_2, n_2)$, one chooses offer a

(b) in (r_1, n_1, r_2, n_2) with a response of "Yes" and $r_1 > r_1^*(y, n_1, r_2, n_2)$, one chooses offer A

(c) in (r_1, n_1, r_2, n_2) with a response of "No" and $r_1 \leq r_1^*(n, n_1, r_2, n_2)$, one chooses offer a

(d) in (r_1, n_1, r_2, n_2) with a response of "No" and $r_1 > r_1^*(n, n_1, r_2, n_2)$, one chooses offer A

For Population 2:

(e) in (r_1, n_1, r_2, n_2) with a response of "Yes" and $r_2 \leq r_2^*(y, r_1, n_1, n_2)$, one chooses offer a

(f) in (r_1, n_1, r_2, n_2) with a response of "Yes" and $r_2 > r_2^*(y, r_1, n_1, n_2)$, one chooses offer A

(g) in (r_1, n_1, r_2, n_2) with a response of "No" and $r_2 \leq r_2^*(n, r_1, n_1, n_2)$, one chooses offer a

(h) in (r_1, n_1, r_2, n_2) with a response of "No" and $r_2 > r_2^*(n, r_1, n_1, n_2)$, one chooses offer A

$r_1^*(n_1, r_2, n_2)$ is non decreasing in n_1, and $r_2^*(r_1, n_1, n_2)$ is also non decreasing in n_2 for both "yes" and "no" responses. We use the experimental results to show some of our findings:

5.4.2 Experimental Results

All the results in this section were run with the same variable values as seen above **Table 5.2**. The strategies seem to imply that there are three regions as r_2 increases. However, "similar offers to all" regions exist. Results of **Table 5.2** suggest that for the region of $r_2 = 1$ to $r_2 = 12$, the action of asking a question does not help decide which offer to extend as the optimal offer strategy is to extend offer a to all. As r_2 increases, there is a change of strategy which is to extend Offer A to customers who say "no". This change of strategy also occurs in the decision for customers who respond with a "yes" at $r_2 = 15$. Then, the strategy is to extend only Offer A to all customers thereafter. However, this is not the case for all instances in (r_1, n_1, r_2, n_2). **Table 5.3** and **Table 5.4** show a counter example where the strategy is to extend Offer A right at the very start to customers. **Table 5.3** starts with those who say "no" and **Table 5.4** to those who say "yes". Customers who respond with a "yes" for **Table 5.3** will be extended Offer A after $r_2 = 4$. In **Table 5.4**, for customers who respond with a "no", Offer A is only extended after $r_2 = 10$.

This suggests that for a small enough r_2, there exists a region where questions are not important. It seems conceivable that there are enough failures to accept that the values for r are low. Hence the optimal offer strategy is to extend only Offer a. Hence for $r_2 \le r_2^*(y/n, n_1, r_2, n_2)$, the retailer chooses Offer a. Conversely, for a large enough r_2, such a region will also exist where questions are not important, but the offer chosen would be of Offer A instead. Hence, there must exist a $r_2^*(r_1, n_1, n_2)$ for responses "yes" and "no", where for $r_2 > r_2^*(y/n, n_1, r_2, n_2)$, Offer A is chosen. This holds true for r_1 also as the optimality equations are identical for both Population 1 and Population 2. Hence the optimal offer strategy described in section 5.4.1 holds. Note that the bold font is used to highlight the decision on the offer and the profit gained from that offer.

$p_1 = 0.65,\ p_2 = 0.60,\ g_1 = g_2 = 10,\ G_1 = G_2 = 25,\ q_1 = 0.2,\ q_2 = 0.4,\ Q_1 = 0.6,\ Q_2 = 0.3,\ o_1 = 0.3,\ o_2 = 0.7$

r_1	n_1	r_2	n_2	MaxProfit	Profit (Yes)	Profit (No)	Decision (Yes)	Decision (No)
26	50	1	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	2	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	3	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	4	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	5	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	6	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	7	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	8	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	9	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	10	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	11	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	12	39	**12.3000**	12.2382	12.3318	***a***	***a***
26	50	13	39	**12.3990**	12.2877	12.4563	***a***	***A***
26	50	14	39	**12.7217**	12.4491	12.8621	***a***	***A***
26	50	15	39	**13.1489**	12.8179	13.3194	***A***	***A***
26	50	16	39	**13.6854**	13.4017	13.8315	***A***	***A***
26	50	17	39	**14.2238**	13.9877	14.3455	***A***	***A***
26	50	18	39	**14.7623**	14.5737	14.8595	***A***	***A***
26	50	19	39	**15.3008**	15.1596	15.3734	***A***	***A***
26	50	20	39	**15.8392**	15.7456	15.8874	***A***	***A***
26	50	21	39	**16.3777**	16.3316	16.4014	***A***	***A***
26	50	22	39	**16.9161**	16.9176	16.9154	***A***	***A***
26	50	23	39	**17.4546**	17.5035	17.4294	***A***	***A***
26	50	24	39	**17.9931**	18.0895	17.9434	***A***	***A***
26	50	25	39	**18.5315**	18.6755	18.4574	***A***	***A***
26	50	26	39	**19.0700**	19.2615	18.9713	***A***	***A***
26	50	27	39	**19.6084**	19.8474	19.4853	***A***	***A***
26	50	28	39	**20.1469**	20.4334	19.9993	***A***	***A***
26	50	29	39	**20.6854**	21.0194	20.5133	***A***	***A***
26	50	30	39	**21.2238**	21.6053	21.0273	***A***	***A***
26	50	31	39	**21.7623**	22.1913	21.5413	***A***	***A***
26	50	32	39	**22.3007**	22.7773	22.0553	***A***	***A***
26	50	33	39	**22.8392**	23.3633	22.5692	***A***	***A***
26	50	34	39	**23.3777**	23.9492	23.0832	***A***	***A***
26	50	35	39	**23.9161**	24.5352	23.5972	***A***	***A***
26	50	36	39	**24.4546**	25.1212	24.1112	***A***	***A***
26	50	37	39	**24.9931**	25.7071	24.6252	***A***	***A***
26	50	38	39	**25.5315**	26.2931	25.1392	***A***	***A***
26	50	39	39	**26.0700**	26.8791	25.6532	***A***	***A***

Table 5.2: Some results from the model

r_1	n_1	r_2	n_2	**MaxProfit**	**Profit (Yes)**	**Profit (No)**	**Decision (Yes)**	**Decision (No)**
50	50	1	15	**12.7800**	12.4782	12.9354	***a***	*A*
50	50	2	15	**13.6200**	12.8982	13.9918	***a***	*A*
50	50	3	15	**14.4606**	13.3186	15.0490	***a***	*A*
50	50	4	15	**15.3728**	13.8616	16.1513	*A*	*A*
50	50	5	15	**16.7500**	15.3603	17.4659	*A*	*A*
50	50	6	15	**18.1500**	16.8838	18.8023	*A*	*A*
50	50	7	15	**19.5500**	18.4073	20.1386	*A*	*A*
50	50	8	15	**20.9500**	19.9309	21.4750	*A*	*A*
50	50	9	15	**22.3500**	21.4544	22.8113	*A*	*A*
50	50	10	15	**23.7500**	22.9779	24.1477	*A*	*A*
50	50	11	15	**25.1500**	24.5014	25.4841	*A*	*A*
50	50	12	15	**26.5500**	26.0250	26.8204	*A*	*A*
50	50	13	15	**27.9500**	27.5485	28.1568	*A*	*A*
50	50	14	15	**29.3500**	29.0720	29.4932	*A*	*A*
50	50	15	15	**30.7500**	30.5956	30.8295	*A*	*A*

Table 5.3: Counter example of offer strategy- Offer *A* for response "No"

r_1	n_1	r_2	n_2	**MaxProfit**	**Profit (Yes)**	**Profit (No)**	**Decision (Yes)**	**Decision (No)**
18	50	1	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	2	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	3	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	4	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	5	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	6	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	7	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	8	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	9	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	10	25	**12.3000**	12.2382	12.3318	*A*	***a***
18	50	11	25	**12.7578**	12.8510	12.7099	*A*	*A*
18	50	12	25	**13.5901**	13.7569	13.5042	*A*	*A*
18	50	13	25	**14.4300**	14.6709	14.3059	*A*	*A*
18	50	14	25	**15.2700**	15.5850	15.1077	*A*	*A*
18	50	15	25	**16.1100**	16.4991	15.9095	*A*	*A*
18	50	16	25	**16.9500**	17.4132	16.7113	*A*	*A*
18	50	17	25	**17.7900**	18.3273	17.5132	*A*	*A*
18	50	18	25	**18.6300**	19.2415	18.3150	*A*	*A*
18	50	19	25	**19.4700**	20.1556	19.1168	*A*	*A*
18	50	20	25	**20.3100**	21.0697	19.9186	*A*	*A*
18	50	21	25	**21.1500**	21.9838	20.7204	*A*	*A*
18	50	22	25	**21.9900**	22.8979	21.5223	*A*	*A*
18	50	23	25	**22.8300**	23.8120	22.3241	*A*	*A*
18	50	24	25	**23.6700**	24.7262	23.1259	*A*	*A*

Table 5.4: Counter example of offer strategy- Offer *A* for response "Yes"

Such conditions of an optimal policy will give possibilities of instances where $r_1^*(n_1, r_2, n_2) = r_2^*(r_1, n_1, n_2)$. But the optimal strategy still holds for this instance as well.

r_1	n_1	r_2	n_2	MaxProfit	Profit (Yes)	Profit (No)	Decision (Yes)	Decision (No)
16	50	1	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	2	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	3	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	4	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	5	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	6	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	7	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	8	20	**12.3000**	12.2382	12.3318	*a*	*a*
16	50	9	20	**12.6165**	12.8117	12.5159	*A*	*A*
16	50	10	20	**13.6207**	13.9049	13.4744	*A*	*A*
16	50	11	20	**14.6700**	15.0468	14.4759	*A*	*A*
16	50	12	20	**15.7200**	16.1894	15.4782	*A*	*A*
16	50	13	20	**16.7700**	17.3320	16.4804	*A*	*A*
16	50	14	20	**17.8200**	18.4747	17.4827	*A*	*A*
16	50	15	20	**18.8700**	19.6173	18.4850	*A*	*A*
16	50	16	20	**19.9200**	20.7600	19.4873	*A*	*A*
16	50	17	20	**20.9700**	21.9026	20.4895	*A*	*A*
16	50	18	20	**22.0200**	23.0453	21.4918	*A*	*A*
16	50	19	20	**23.0700**	24.1879	22.4941	*A*	*A*
16	50	20	20	**24.1200**	25.3306	23.4963	*A*	*A*

Table 5.5: Same offers for same responses to questions

We cannot prove these results using the technique of Theorem 3.1 because the equation of s when used in this case would split into s_1 for Population 1 and s_2 for Population 2, where $s = s_1 + s_2$. The proof will not hold as the use of the s will not be consistent through the equations and the effect of increment of r not as obvious as it is.

5.5 Known Probabilities of Both Offer *a* and Offer *A*, Incomplete Knowledge of Origin of Customer - Two Questions

This case is similar to **section 5.3** where the retailer has complete knowledge of the acceptance probabilities of both offers for population 1 and population 2. Hence, there is no learning performed here and that means that the retailer makes the same decision (same offer) to all the applicants irrespective of how many more are to be considered. We introduce question q and question Q as the two questions where only one is selected to be asked. The optimality equations are categorized according to question and response.

If response to the question q is "yes" (y):

$$v(y) = \max\begin{cases} \left(\dfrac{q_1 o_1}{q_1 o_1 + q_2 o_2}\right) p_1 g_1 + \left(\dfrac{q_2 o_2}{q_1 o_1 + q_2 o_2}\right) p_2 g_2, \\ \left(\dfrac{q_1 o_1}{q_1 o_1 + q_2 o_2}\right) P_1 G_1 + \left(\dfrac{q_2 o_2}{q_1 o_1 + q_2 o_2}\right) P_2 G_2. \end{cases}$$

If the response to the first question is a "no":

$$v(n) = \max\begin{cases} \left(\dfrac{(1-q_1)o_1}{(1-q_1)o_1 + (1-q_2)o_2}\right) p_1 g_1 + \left(\dfrac{(1-q_2)o_2}{(1-q_1)o_1 + (1-q_2)o_2}\right) p_2 g_2, \\ \left(\dfrac{(1-q_1)o_1}{(1-q_1)o_1 + (1-q_2)o_2}\right) P_1 G_1 + \left(\dfrac{(1-q_2)o_2}{(1-q_1)o_1 + (1-q_2)o_2}\right) P_2 G_2. \end{cases}$$

The profit from an applicant if answering question q, $v(q, r_1, n_1, r_2, n_2)$ is as follows:

$$v(q, r_1, n_1, r_2, n_2) = (q_1 o_1 + q_2 o_2)v(y) + ((1-q_1)o_1 + (1-q_2)o_2)v(n).$$

Note that for both instances, offer a is extended if the following condition is fulfilled:

(i) $$\left(\frac{q_1}{q_2}\right)\left(\frac{o_1}{1-o_1}\right) > \frac{P_2 G_2 - p_2 g_2}{p_1 g_1 - P_1 G_1}.$$

(ii) $$\left(\frac{1-q_1}{1-q_2}\right)\left(\frac{o_1}{1-o_1}\right) > \frac{P_2 G_2 - p_2 g_2}{p_1 g_1 - P_1 G_1}.$$

Assume that

Q_1 = Probability of Population 1 saying "Yes" to the question Q

Q_2 = Probability of Population 2 saying "Yes" to the question Q

If response to the question Q is "Yes" (Y):

$$v(Y) = \max\begin{cases} \left(\dfrac{Q_1 o_1}{Q_1 o_1 + Q_2 o_2}\right) p_1 g_1 + \left(\dfrac{Q_2 o_2}{Q_1 o_1 + Q_2 o_2}\right) p_2 g_2, \\ \left(\dfrac{Q_1 o_1}{Q_1 o_1 + Q_2 o_2}\right) P_1 G_1 + \left(\dfrac{Q_2 o_2}{Q_1 o_1 + Q_2 o_2}\right) P_2 G_2. \end{cases}$$

If the response to the question Q is a "No":

$$v_i(N) = \max\begin{cases} \left(\dfrac{(1-Q_i) o_i}{(1-Q_1) o_1 + (1-Q_2) o_2}\right) p_i g_i, \\ \left(\dfrac{(1-Q_i) o_i}{(1-Q_1) o_1 + (1-Q_2) o_2}\right) P_i G_i. \end{cases}$$

Hence the profit from the applicant before answering question Q:

$$v(Q, r_1, n_1, r_2, n_2) = (Q_1 o_1 + Q_2 o_2) v(y) + ((1-Q_1) o_1 + (1-Q_2) o_2) v(n).$$

The condition for offer a is the same as when the question Q is asked:

(iii) $\left(\dfrac{Q_1}{Q_2}\right)\left(\dfrac{o_1}{1-o_1}\right) > \dfrac{P_2 G_2 - p_2 g_2}{p_1 g_1 - P_1 G_1}.$

(iv) $\left(\dfrac{1-Q_1}{1-Q_2}\right)\left(\dfrac{o_1}{1-o_1}\right) > \dfrac{P_2 G_2 - p_2 g_2}{p_1 g_1 - P_1 G_1}.$

We have to decide which question to ask first and use the question that will give the highest expected profit. Hence the expected maximum profit of the applicant before answering a chosen question is:

$$v(r_1, n_1, r_2, n_2) = \max\begin{cases} v(q, r_1, n_1, r_2, n_2), \\ v(Q, r_1, n_1, r_2, n_2). \end{cases}$$

The decision on the offer strategy for this case is a ratio of $\frac{Q_1}{Q_2}, \frac{1-Q_1}{1-Q_2}, \frac{q_1}{q_2}, \frac{1-q_1}{1-q_2}$ and $\frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$. So the optimal offer strategy can be summarised as follows:

(a) $$\frac{q_1}{q_2} > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer a for a "yes" response to question q.

(b) $$\frac{q_1}{q_2} \leq \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer A for a "yes" response to question q.

(c) $$\frac{1-q_1}{1-q_2} > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer a for a "no" response to question q.

(d) $$\frac{1-q_1}{1-q_2} \leq \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer A for a "no" response to question q.

(e) $$\frac{Q_1}{Q_2} > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer a when there is a "yes" response to question Q.

(f) $$\frac{Q_1}{Q_2} \leq \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer A when there is a "yes" response to question Q.

(g) $$\frac{1-Q_1}{1-Q_2} > \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer a when there is a "no" response to question Q.

(h) $$\frac{1-Q_1}{1-Q_2} \leq \frac{P_2G_2 - p_2g_2}{p_1g_1 - P_1G_1}\left(\frac{1-o_1}{o_1}\right)$$

Extend Offer A when there is a "no" response to question Q.

5.5.1 Experimental Results

We present an example of experimental results from the model with the following values:

(Note the bold font highlights the decision made on the question to be asked, the corresponding offer and the profit gained from the offer extended).

$p_1 = 0.65$, $p_2 = 0.60$, $P_1 = 0.3$, $P_2 = 0.5$, $g_1 = g_2 = 10$, $G_1 = G_2 = 25$, $q_1 = 0.2$, $_2 = 0.4$, $Q_1 = 0.6$, $Q_2 = 0.3$, $o_1 = 0.3$, $o_2 = 0.7$

r_1	n_1	r_2	n_2	**Profit(*q*)**	**ProfitY*q***	**ProfitN*q***	**DeciY*q***	**DeciN*q***	**Profit(*Q*)**	**ProfitY*Q***	**ProfitN*Q***	**MaxProfit**	**DeciY*Q***	**DeciN*Q***
22	32	1	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	2	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	3	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	4	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	5	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	6	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	7	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	8	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	9	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	10	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	11	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	12	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	13	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	14	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	15	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	16	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	17	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	18	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	19	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	20	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	21	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	22	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	23	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*
22	32	24	24	11.0000	11.6176	10.6818	*A*	*A*	11.0000	10.1923	11.5164	**11.0000**	*A*	*A*

Table 5.6: Example of offer strategy if both P_i and p_i is known

If all the acceptance probabilities for all the two offers are known, hence only one offer will be extended as found in the results. Hence no learning has been performed and the model only selects the offer which gives the highest maximum expected future profit.

5.6 Known Probability of Acceptance for Offer a (p_i) but Unknown Probability of Acceptance for Offer A (P_i), Incomplete Knowledge of Origin of Customer – Two Questions

We maintain all the variables used in **Section 5.4** and **Section 5.5** for the case of having complete information on the acceptance probability of Offer a, but not of Offer A, and having two questions to ask the applicant. The optimality equations for this case are as follows:

If response to question q is "yes" (y):

$$v(y, r_1, n_1, r_2, n_2) = \max\begin{cases} \frac{q_1 o_1}{q_1 o_1 + q_2 o_2} p_1 g_1 + \frac{q_2 o_2}{q_1 o_1 + q_2 o_2} p_2 g_2 + \beta v(r_1, n_1, r_2, n_2), \\ \frac{q_1 o_1}{q_1 o_1 + q_2 o_2} p_1 \frac{r_1}{n_1} G_1 + \frac{q_2 o_2}{q_1 o_1 + q_2 o_2} p_2 \frac{r_2}{n_2} G_2 + \\ \beta(\frac{q_1 o_1}{q_1 o_1 + q_2 o_2}\{p_1 \frac{r_1}{n_1} v(r_1+1, n_1+1, r_2, n_2) + p_1\left(1-\frac{r_1}{n_1}\right) v(r_1, n_1+1, r_2, n_2) + \\ (1-p_1) v(r_1, n_1, r_2, n_2)\} + \frac{q_2 o_2}{q_1 o_1 + q_2 o_2}\{p_2 \frac{r_2}{n_2} v(r_1, n_1, r_2+1, n_2+1) + \\ p_2\left(1-\frac{r_2}{n_2}\right) v(r_1, n_1, r_2, n_2+1) + (1-p_2) v(r_1, n_1, r_2, n_2)\}). \end{cases}$$

If the response to question q is a "no":

$$v(n, r_1, n_1, r_2, n_2) = \max\begin{cases} \frac{(1-q_1)o_1}{(1-q_1)o_1 + (1-q_2)o_2} p_1 g_1 + \frac{(1-q_2)o_2}{(1-q_1)o_1 + (1-q_2)o_2} p_2 g_2 + \beta v(r_1, n_1, r_2, n_2), \\ \frac{(1-q_1)o_1}{(1-q_1)o_1 + (1-q_2)o_2} p_1 \frac{r_1}{n_1} G_1 + \frac{(1-q_2)o_2}{(1-q_1)o_1 + (1-q_2)o_2} p_2 \frac{r_2}{n_2} G_2 + \\ \beta(\frac{(1-q_1)o_1}{(1-q_1)o_1 + (1-q_2)o_2}\{p_1 \frac{r_1}{n_1} v(r_1+1, n_1+1, r_2, n_2) + p_1\left(1-\frac{r_1}{n_1}\right) v(r_1, n_1+1, r_2, n_2) + \\ (1-p_1) v(r_1, n_1, r_2, n_2)\} + \frac{(1-q_2)o_2}{(1-q_1)o_1 + (1-q_2)o_2}\{p_2 \frac{r_2}{n_2} v(r_1, n_1, r_2+1, n_2+1) + \\ p_2\left(1-\frac{r_2}{n_2}\right) v(r_1, n_1, r_2, n_2+1) + (1-p_2) v(r_1, n_1, r_2, n_2)\}). \end{cases}$$

We define the profit from the applicant before answering question q as:

$$v(q,r_1,n_1,r_2,n_2)=(q_1o_1+q_2o_2)v(y,r_1,n_1,r_2,n_2)+((1-q_1)o_1+(1-q_2)o_2)v(n,r_1,n_1,r_2,n_2).$$

The optimality equations for question Q are as follows:

If response to question Q is "Yes" (Y):

$$v(Y,r_1,n_1,r_2,n_2)=\max\begin{cases}\frac{Q_1o_1}{Q_1o_1+Q_2o_2}p_1g_1+\frac{Q_2o_2}{Q_1o_1+Q_2o_2}p_2g_2+\beta v(r_1,n_1,r_2,n_2),\\ \frac{Q_1o_1}{Q_1o_1+Q_2o_2}p_1\frac{r_1}{n_1}G_1+\frac{Q_2o_2}{Q_1o_1+Q_2o_2}p_2\frac{r_2}{n_2}G_2+\\ \beta(\frac{Q_1o_1}{Q_1o_1+Q_2o_2}\{p_1\frac{r_1}{n_1}v(r_1+1,n_1+1,r_2,n_2)+p_1\left(1-\frac{r_1}{n_1}\right)v(r_1,n_1+1,r_2,n_2)+\\ (1-p_1)v(r_1,n_1,r_2,n_2)\}+\frac{Q_2o_2}{Q_1o_1+Q_2o_2}\{p_2\frac{r_2}{n_2}v(r_1,n_1,r_2+1,n_2+1)+\\ p_2\left(1-\frac{r_2}{n_2}\right)v(r_1,n_1,r_2,n_2+1)+(1-p_2)v(r_1,n_1,r_2,n_2)\}).\end{cases}$$

If the response to question Q is a "No":

$$v(N,r_1,n_1,r_2,n_2)=\max\begin{cases}\frac{(1-Q_1)o_1}{(1-Q_1)o_1+(1-Q_2)o_2}p_1g_1+\frac{(1-Q_2)o_2}{(1-Q_1)o_1+(1-Q_2)o_2}p_2g_2+\beta v(r_1,n_1,r_2,n_2),\\ \frac{(1-Q_1)o_1}{(1-Q_1)o_1+(1-Q_2)o_2}p_1\frac{r_1}{n_1}G_1+\frac{(1-Q_2)o_2}{(1-Q_1)o_1+(1-Q_2)o_2}p_2\frac{r_2}{n_2}G_2+\\ \beta(\frac{(1-Q_1)o_1}{(1-Q_1)o_1+(1-Q_2)o_2}\{p_1\frac{r_1}{n_1}v(r_1+1,n_1+1,r_2,n_2)+p_1\left(1-\frac{r_1}{n_1}\right)v(r_1,n_1+1,r_2,n_2)+\\ (1-p_1)v(r_1,n_1,r_2,n_2)\}+\frac{(1-Q_2)o_2}{(1-Q_1)o_1+(1-Q_2)o_2}\{p_2\frac{r_2}{n_2}v(r_1,n_1,r_2+1,n_2+1)+\\ p_2\left(1-\frac{r_2}{n_2}\right)v(r_1,n_1,r_2,n_2+1)+(1-p_2)v(r_1,n_1,r_2,n_2)\}).\end{cases}$$

Hence the profit from the applicant before answering question Q is:

$$v(Q,r_1,n_1,r_2,n_2)=(Q_1o_1+Q_2o_2)v(Y,r_1,n_1,r_2,n_2)+((1-Q_1)o_1+(1-Q_2)o_2)v(N,r_1,n_1,r_2,n_2).$$

We use the same strategy as in Section 5.6 to calculate the expected maximum future profit of the next customer and select the question to ask:

$$v(r_1, n_1, r_2, n_2) = \max\begin{cases} v(q, r_1, n_1, r_2, n_2), \\ v(Q, r_1, n_1, r_2, n_2). \end{cases}$$

5.6.1 Experimental Results

In this subsection, we present some experimental results of the case in Section 5.6 and show some proofs of the results later. We maintain Population 1 and Population 2 and the variables we used earlier in this chapter. But, we consider now that we can ask two questions. The objective is to decide which of the two questions to ask first. We decide this by choosing the question which maximises the expected future profit of the customer. The following tables show the experimental results. (Note the bold font highlights the decision made on the question to be asked, the corresponding offer and the profit gained from the offer extended).

$p_1 = 0.65$, $p_2 = 0.60$, $g_1 = g_2 = 10$, $G_1 = G_2 = 25$, $q_1 = 0.2$, $q_2 = 0.4$, $Q_1 = 0.6$, $Q_2 = 0.3$, $o_1 = 0.3$, $o_2 = 0.7$

r_1	n_1	r_2	n_2	**Profit(*q*)**	**ProfitY*q***	**ProfitN*q***	**DeciY*q***	**DeciN*q***	**Profit(*Q*)**	**ProfitY*Q***	**ProfitN*Q***	**MaxProfit**	**DeciY*Q***	**DeciN*Q***
45	45	1	24	13.0087	12.8645	13.083	*a*	*A*	**13.5525**	14.6128	12.8746	13.5525	*A*	***a***
45	45	2	24	13.4025	12.9957	13.612	*a*	*A*	**13.815**	15.0806	13.0058	13.815	*A*	***a***
45	45	3	24	13.7962	13.127	14.141	*a*	*A*	**14.0775**	15.5483	13.1371	14.0775	*A*	***a***
45	45	4	24	14.19	13.2582	14.6701	*a*	*A*	**14.34**	16.0162	13.2684	14.34	*A*	***a***
45	45	5	24	14.5898	13.3915	15.2072	*a*	*A*	**14.6066**	16.4912	13.4016	14.6066	*A*	***a***
45	45	6	24	**15.0988**	13.6376	15.8515	***a***	*A*	15.0567	17.0742	13.7668	15.0988	*A*	*A*
45	45	7	24	15.8752	14.4083	16.6309	*A*	*A*	15.8752	17.7934	14.6488	15.8752	*A*	*A*
45	45	8	24	16.75	15.3603	17.4659	*A*	*A*	16.75	18.5673	15.5881	16.75	*A*	*A*
45	45	9	24	17.625	16.3125	18.3011	*A*	*A*	17.625	19.3413	16.5276	17.625	*A*	*A*
45	45	10	24	18.5	17.2647	19.1363	*A*	*A*	18.5	20.1154	17.4672	18.5	*A*	*A*
45	45	11	24	19.375	18.2169	19.9716	*A*	*A*	19.375	20.8894	18.4067	19.375	*A*	*A*
45	45	12	24	20.25	19.1691	20.8068	*A*	*A*	20.25	21.6634	19.3463	20.25	*A*	*A*
45	45	13	24	21.125	20.1213	21.642	*A*	*A*	21.125	22.4375	20.2858	21.125	*A*	*A*
45	45	14	24	22	21.0735	22.4772	*A*	*A*	22	23.2115	21.2254	22	*A*	*A*
45	45	15	24	22.875	22.0257	23.3125	*A*	*A*	22.875	23.9855	22.1649	22.875	*A*	*A*
45	45	16	24	23.75	22.9779	24.1477	*A*	*A*	23.75	24.7596	23.1045	23.75	*A*	*A*
45	45	17	24	24.625	23.9301	24.9829	*A*	*A*	24.625	25.5336	24.044	24.625	*A*	*A*
45	45	18	24	25.5	24.8823	25.8181	*A*	*A*	25.5	26.3076	24.9836	25.5	*A*	*A*
45	45	19	24	26.375	25.8345	26.6534	*A*	*A*	26.375	27.0817	25.9231	26.375	*A*	*A*
45	45	20	24	27.2499	26.7867	27.4886	*A*	*A*	27.2499	27.8557	26.8627	27.2499	*A*	*A*
45	45	21	24	28.1249	27.7389	28.3238	*A*	*A*	28.1249	28.6298	27.8022	28.1249	*A*	*A*
45	45	22	24	28.9999	28.6911	29.159	*A*	*A*	28.9999	29.4038	28.7417	28.9999	*A*	*A*
45	45	23	24	29.8749	29.6433	29.9943	*A*	*A*	29.8749	30.1778	29.6813	29.8749	*A*	*A*
45	45	24	24	30.7499	30.5955	30.8295	*A*	*A*	30.7499	30.9519	30.6208	30.7499	*A*	*A*

Table 5.7: An example of results for two questions

Profit(q) is the profit from the applicant before answering question q:
$v(q, r_1, n_1, r_2, n_2) = (q_1 o_1 + q_2 o_2) v(y, r_1, n_1, r_2, n_2) + ((1-q_1) o_1 + (1-q_2) o_2) v(n, r_1, n_1, r_2, n_2)$.

Profit yq is $v(y, r_1, n_1, r_2, n_2)$ and Profit nq is $v(n, r_1, n_1, r_2, n_2)$ while Deci yq and Deci nq are the decisions on the offers for question q for responses "yes" or "no".

Hence Profit(Q) is the expected maximum future profit from the next customer if question Q is chosen which is defined as:
$v(Q, r_1, n_1, r_2, n_2) = (Q_1 o_1 + Q_2 o_2) v(Y, r_1, n_1, r_2, n_2) + ((1-Q_1) o_1 + (1-Q_2) o_2) v(N, r_1, n_1, r_2, n_2)$
where ProfitYQ is $v(Y, r_1, n_1, r_2, n_2)$ and ProfitNQ is $v(N, r_1, n_1, r_2, n_2)$ while DeciYQ and DeciNQ are the decisions on the offers for question Q given the response to the question is "Yes" or "No". ProfitNew is defined as the maximum expected future profit of the next customer not knowing what responses will be given to the question being asked and is defined as:

$$v(r_1, n_1, r_2, n_2) = \max \begin{cases} v(q, r_1, n_1, r_2, n_2), \\ v(Q, r_1, n_1, r_2, n_2). \end{cases}$$

The results show that given there are two questions, there exists three interesting regions that make up the offer strategy. The first being that for small enough values of r_i, it doesn't matter which question is selected, the retailer extends Offer a to all the customers. Hence there could be enough rejections of offers which causes the r_i to be low. Conversely, for the case of $r_i = n_i$, Offer A is offered. Hence there also exist an area that is bounded by $r_i = n_i$ that the retailer will extend only Offer A to all the customers. Therefore, we can summarise this as follows (for reasons of simplicity, we assume that r_1 and n_1 is constant):

(i) for (r_1, n_1, r_2, n_2), for $r_2 \le r_2^*(r_1, n_1, n_2)$ extend Offer a.

(ii) for (r_1, n_1, r_2, n_2), if $r_2 > r_2^*(r_1, n_1, n_2)$ and $r_1 > r_1^*(n_1, r_2, n_2)$ extend Offer A.

At the beginning, the base offer which is Offer a is extended to all customers. We know that the customer will accept the base offer. Hence which question is asked does not really matter. The second region is where the questions play an important role in maximising the expected future profit. The question is chosen based on the objective of maximising the future profit of the customer, where Offer A will be preferred over Offer a. The model will learn and keep updating itself based on the acceptance behaviour of the future customers and decide when it should start extending Offer A. When the model becomes confident enough that a group will accept Offer A it will extend Offer A. Hence, the questions will not be relevant anymore. There will also exist a state where r_1 is really small compared with n_2, then one might only Offer a in state (r_1, n_1, n_2, n_2). For example, given $o_1 = 0.9$ and $o_2 = 0.1$ (rest of the variables unchanged from subsection 5.4.2) with $r_1 = 4$ and $n_2 = 43$, the decision is to extend Offer a.

But at the state where the questions do help on deciding the offers, the questions chosen can start with, for example question Q and at some point switch to question q, see **Table 5.9**. In **Table 5.10**, where the offer strategy begins with Offer a, note that the profit is constant as long as only Offer a is extended. As the model starts extending Offer A, the value will start to increase. Hence once $r_2 > r_2^*(r_1, n_1, n_2)$, Offer A is chosen. Hence the offer strategy is (i) and (ii) holds.

r_1	n_1	r_2	n_2	**Profit(q)**	**Profity$q$**	**Profitnq**	**Deci yq**	**Deci nq**	**Profit(Q)**	**ProfitYQ**	**ProfitNQ**	**MaxProfit**	**DeciYQ**	**DeciNQ**
45	45	1	11	13.4741	13.0196	13.7082	*a*	*A*	**13.8627**	15.1656	13.0297	13.8627	*A*	***a***
45	45	2	11	14.341	13.3085	14.8729	*a*	*A*	**14.4407**	16.1955	13.3187	14.4407	*A*	***a***
45	45	3	11	15.5013	14.0014	16.274	*A*	*A*	15.5013	17.4627	14.2473	15.5013	*A*	*A*
45	45	4	11	17.3864	16.0528	18.0733	*A*	*A*	17.3864	19.1302	16.2714	17.3864	*A*	*A*
45	45	5	11	19.2954	18.1303	19.8956	*A*	*A*	19.2954	20.819	18.3213	19.2954	*A*	*A*
45	45	6	11	21.2045	20.2078	21.7179	*A*	*A*	21.2045	22.5078	20.3712	21.2045	*A*	*A*
45	45	7	11	23.1136	22.2854	23.5402	*A*	*A*	23.1136	24.1966	22.4212	23.1136	*A*	*A*
45	45	8	11	25.0227	24.3629	25.3626	*A*	*A*	25.0227	25.8854	24.4711	25.0227	*A*	*A*
45	45	9	11	26.9318	26.4405	27.1849	*A*	*A*	26.9318	27.5743	26.521	26.9318	*A*	*A*
45	45	10	11	28.8409	28.518	29.0072	*A*	*A*	28.8409	29.2631	28.5709	28.8409	*A*	*A*
45	45	11	11	30.7499	30.5955	30.8295	*A*	*A*	30.7499	30.9519	30.6208	30.7499	*A*	*A*

Table 5.8: Only asking question Q as part of the offer strategy

r_1	n_1	r_2	n_2	**Profit(q)**	**Profit yq**	**Profit nq**	**Deci yq**	**Deci nq**	**Profit(Q)**	**ProfitYQ**	**ProfitNQ**	**MaxProfit**	**DeciYQ**	**DeciNQ**
45	45	1	13	13.3419	12.9755	13.5306	*a*	*A*	**13.7746**	15.0086	12.9856	13.7746	***A***	***a***
45	45	2	13	14.0693	13.218	14.5079	*a*	*A*	**14.2595**	15.8728	13.2281	14.2595	***A***	***a***
45	45	3	13	**14.8605**	13.5185	15.5518	***a***	***A***	14.8053	16.8023	13.5286	14.8605	*A*	*a*
45	45	4	13	16.2125	14.7754	16.9528	***A***	*A*	16.2125	18.0918	15.011	16.2125	*A*	***A***
45	45	5	13	17.8269	16.5322	18.4938	***A***	*A*	17.8269	19.5199	16.7445	17.8269	*A*	***A***
45	45	6	13	19.4423	18.2901	20.0358	***A***	*A*	19.4423	20.9489	18.479	19.4423	*A*	***A***
45	45	7	13	21.0577	20.048	21.5778	***A***	*A*	21.0577	22.3779	20.2135	21.0577	*A*	***A***
45	45	8	13	22.673	21.806	23.1197	***A***	*A*	22.673	23.8069	21.9481	22.673	*A*	***A***
45	45	9	13	24.2884	23.5639	24.6617	***A***	*A*	24.2884	25.2359	23.6826	24.2884	*A*	***A***
45	45	10	13	25.9038	25.3218	26.2036	***A***	*A*	25.9038	26.6649	25.4172	25.9038	*A*	***A***
45	45	11	13	27.5192	27.0797	27.7456	***A***	*A*	27.5192	28.0939	27.1517	27.5192	*A*	***A***
45	45	12	13	29.1346	28.8376	29.2875	***A***	*A*	29.1346	29.5229	28.8863	29.1346	*A*	***A***
45	45	13	13	30.7499	30.5955	30.8295	***A***	*A*	30.7499	30.9519	30.6208	30.7499	*A*	***A***

Table 5.9: Example of asking question Q and then q as part of the offer strategy

r_1	n_1	r_2	n_2	Profit(q)	Profit yq	Profit nq	Deci yq	Deci nq	Profit(Q)	ProfitYQ	ProfitNQ	MaxProfit	DeciYQ	DeciNQ
1	27	1	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	2	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	3	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	4	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	5	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	6	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	7	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	8	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	9	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	10	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	11	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	12	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	13	28	12.3	12.2382	12.3318	*a*	*a*	12.3	12.3807	12.2483	12.3	*a*	*a*
1	27	14	28	**12.436**	12.5062	12.3998	*A*	*a*	12.399	12.4487	12.3664	12.436	*a*	*A*
1	27	15	28	12.807	13.1648	12.622	*A*	*a*	**12.8804**	12.6709	13.0143	12.8804	*a*	*A*
1	27	16	28	13.218	13.8672	12.884	*A*	*a*	**13.4044**	12.933	13.7059	13.4044	*a*	*A*
1	27	17	28	13.631	14.5709	13.1465	*A*	*a*	**13.9294**	13.1955	14.3987	13.9294	*a*	*A*
1	27	18	28	14.158	15.2746	13.5824	*A*	*A*	**14.454**	13.458	15.0915	14.4544	*a*	*A*
1	27	19	28	14.795	15.9783	14.1858	*A*	*A*	**14.979**	13.7205	15.7843	14.9794	*a*	*A*
1	27	20	28	15.433	16.6824	14.7898	*A*	*A*	**15.5048**	13.9832	16.4776	15.5048	*a*	*A*
1	27	21	28	16.1178	17.4337	15.4399	*A*	*A*	16.1178	14.397	17.218	16.1178	*A*	*A*
1	27	22	28	16.8612	18.2428	16.1495	*A*	*A*	16.8612	15.0546	18.0163	16.8612	*A*	*A*
1	27	23	28	17.6111	19.0588	16.8653	*A*	*A*	17.6111	15.7179	18.8215	17.6111	*A*	*A*
1	27	24	28	18.3611	19.875	17.5812	*A*	*A*	18.3611	16.3814	19.6268	18.3611	*A*	*A*
1	27	25	28	19.1111	20.6911	18.2971	*A*	*A*	19.1111	17.0448	20.4321	19.1111	*A*	*A*
1	27	26	28	19.8611	21.5073	19.013	*A*	*A*	19.8611	17.7083	21.2374	19.8611	*A*	*A*
1	27	27	28	20.6111	22.3235	19.7289	*A*	*A*	20.6111	18.3718	22.0428	20.6111	*A*	*A*
1	27	28	28	21.3611	23.1397	20.4448	*A*	*A*	21.3611	19.0352	22.8481	21.3611	*A*	*A*

Table 5.10: Example where are three regions exist, and the question asked first is question q then a switch to Q

5.7 Conclusion

In this chapter, we have explored the different cases of predicting the offer strategy of two offers; Offer a and Offer A. We have looked at the strategies given complete (incomplete) knowledge of the acceptance probabilities and also complete (incomplete) knowledge of the origin of the customers. We also look asking one question or two questions to help decide on the offers to extend to the customers. We have managed to formulate an offer strategy for both of these cases as well.

Chapter 6

6.0 Conclusions

This concluding chapter summarises the main results from this piece of research and includes a section on future work. The aim of this research was to use adaptive learning in credit scoring to estimate the acceptance probability distribution for a financial product.

So in Chapter 3, we built an acceptance model based on dynamic programming with Bayesian elements to include past actions in the optimality equations. We first built the model to predict the best offer to make to the next customer, for two different variants of a credit card. Experimental results and mathematical induction were used to prove the monotonicity characteristics of the maximum expected future profit of the customer. This profit function is updated as the customer decides which of the two offers to accept. Mathematical proof was also provided to validate the optimal offer strategy. We also extended the model to accommodate N variants of the offer and proved the monotonicity characteristics of the maximum expected future profit and the optimal offer strategy

Next, we looked at a model to help decide on which offer (with high acceptance probability) to make to a customer with certain characteristics based on past information. In Chapter 4, we built TAROT to select a question (questions) to ask customers to match the responses with an offer with high acceptance. We show step by step how to build TAROT using Enterprise Miner 4.3 from SAS 9.1.3. If the retailer wants to ask m number of questions, hence the TAROT tree built will have m levels of applicant characteristics followed by n level of offers. The questions are based on the applicant characteristic splits. The applicant-offer sequence of TAROT based on the classification power of CART succeeds in deciding on the questions to identify the most profitable and highly accepted offers.

Finally in Chapter 5, we utilise the profit updates to select between two questions, q and Q. The questions in turn help decide on the optimal offer strategy between two offers. We provide the proofs and experimental results for different cases, namely when the acceptance probability of both offers is known, and when only one of the acceptance probabilities of the offers is known. The optimal offer strategy is then obtained.

6.1 Future Research

6.1.1 Acceptance Model

The acceptance model introduced in Chapter 3 is based on dynamic programming with Bayesian elements to allow the inclusion of past actions into the optimality equations. The results we presented in Chapter 3 where there were N variants of a product but they were different choice of one function like the interest rate charged. An interesting experiment would be dealing with more than one dimensional offer characteristics, which is offer (i, j) where this is offer of level i in a offer characteristic 1 and j in a offer characteristic 2. This would correspond to a problem where one is deciding on both interest rate to charge and credit limit to offer.

6.1.2 TAROT (Top Applicant characteristics Remainder Offer characteristics Tree)

The TAROT that we built in Chapter 4 is used to help decide which offer to make to customers based on their responses to questions on selected applicant characteristics. We used the powerful classification powers of a CART which is incorporated within TAROT to identify the significant applicant and offer characteristics. We also used the acceptance rates and later for the bootstrapped data the acceptance probabilities for the variables to decide which offers to make. A worthwhile future extension will be using TAROT to identify another offer for purpose of cross-selling. Also, it would be interesting to look at the viability of updating the data on which the TAROT is based on. TAROT is built entirely on past data, but the question will be for how long the data will accurately portray the trend of the population. And if the decision is to update the data, how long should one wait before this "recalibration" exercise needs to be performed?

6.1.3 Choosing Questions

In Chapter 5, we combined the acceptance model from Chapter 3 and the questions selection process from TAROT (Chapter 4) to help decide on the selection of questions to ask so as to match the responses with the offers with high acceptance rates. Instead of using past data to help decide on the questions (like TAROT), we used the updated maximum expected future profit function introduced in Chapter 3 to choose which question to ask. In reality, updating after every offer will not be realistic and yet the acceptance rates will be influenced by new market or economic conditions. Hence there are two areas for future research here. The first is when to do the block calculations for the population at hand; should it be daily, weekly, monthly or quarterly? The second is to find out how to treat the effects of the new conditions of the market or economy for the model. For example, does one just decide to ignore the old data or discount the old data and include it in the calculations? These would be areas that could be looked into for further research.

Appendix

Page 1 of Fantasy Student Account Website

Your Chance To Win Money!!

University of Southampton

£50 WIN WIN WIN £50

WIN

Fantasy Account Chooser

You could win £50 cash and maybe even help students get a better deal on their bank accounts in future years for only 5 minutes work

All You have to do is fill in the following form as if you were applying for a student bank account

You will then be offered an account based on that data. Select whether you would accept this account in the real world and then fill out the feedback form.

All details are purely for a research project and no companies will have access to you data.

One lucky person chosen at random will be sent £50 for taking part - It could be you.

Click to Play

Page 2 of Fantasy Student Account Website

Fantasy Account Application Form

Please take the time to fill out this application form in order for us to set up your account.

*Fields Marked with * must be filled in*

Personal Details:

** Forename*

** Surname*

Date of Birth: Day: 1 ** Month* Jan *Year*

Sex: *Male* *Female*

Status: *Married* *Single* *Divorced* *Widowed/Separated* *Other*

No. Of Dependent Children:

** E-mail Address (So we can contact you about your prize):*

name@internet.address

Financial Details

How many credit cards/store cards do you hold: 0

Income & Expenditure Details

Which of these do you recieve:

Wage *Student Loan/ Grant* *Parental Contribution*

Other please specify:

College/University Details

Establishment: University

Name of University and Campus/College:

Course Type:

Law *Economics* *Education* *Medicine*

Engineering *Sciences* *Arts* *Management*

Languages *Nursing* *Computer Sciences* *Other*

Please check any of the following that you have an interest in:

Sport *Travel* *Music* *Clubbing*

Cinema *Cars* *DIY* *Gardening* *Beer* *Country and Western Music*

Submit Form | Reset Form

Page 3 of Fantasy Student Account

Select Your Account

Based on the data you submitted in your application form, we have selected the following bank accounts to suit your needs.
Please take the time to look at what the accounts offer and then decide whether you would take either one of them or not.
Please remember this is all just research and we are not going to send you any junk mail, so answer as honestly as possible and you could be in with a chance of winning £50

Hello Y! You're looking lovely today.

You are user number 956 and the date is Tuesday the 28 of March, just a little bit of trivia there for you, hope it makes you happy. Which of these accounts if any suit your needs?

Please check the box of the account you prefer

Accept This Maxi Account?

○ Maxi

Maxi Account					
Credit Card Offer/Limit	**Overdraft Limit**	**Commision -Free Travel Money**	**Insurance Deal (s)**	**Introductory Offer**	**Interest when in Credit**
£500 Credit limit	£1250	No	15-20% Discount on insuring Instruments\CD's+records\Decks	5 Free CD's from selection of our wicked stockpile inc. the best of Militant Jungle/Drum 'n' Bass/ Hardcore/ Techno/ Deep Trance/ House etc...	0%

Or do you prefer this Million account?

○ Million

Million Account					
Credit Card Offer/Limit	**Overdraft Limit**	**Commision-Free Travel Money**	**Insurance Deal (s)**	**Introductory Offer**	**Interest when in Credit**
No Credit Card	£1700	Yes	None	None	0%

Or perhaps you think neither of these accounts offers the deal for you in which case, please check this box

○ Neither

Submit Form

References

1 Ahn B. S., Park K. S., Han C. H. and Kim J. K. (2000). Multi-attribute decision aid under incomplete information and hierarchical structure. *European Journal of Operational Research.* **125:** 431-439.

2 Al-Khayyal F., Griffin P. M. and Smith N. R. (2001). Solution of a large-scale two-stage decision and scheduling problem using decomposition. *European Journal of Operational Research.* **132:** 453-465.

3 Baesens B. (2003) *Developing intelligent systems for credit scoring using machine learning techniques*, Department of Applied Economic Sciences, Katholieke Universiteit Leuven. pp: 262.

4 Baesens B., Gestel T. V., Viaene S., Stepanova M., Suykens J. and Vanthienen J. (2003). Benchmarking state-of-the-art classification algorithms for credit scoring. *Journal of the Operational Research Society.* **54:** 627-635.

5 Baesens B., Viaene S., Van den Poel D., Vanthienen J. and Dedene G. (2002). Bayesian neural network learning for repeat purchase modelling in direct marketing. *European Journal of Operational Research.* **138:** 191-211.

6 Bellman R. (1957) *Dynamic Programming,* Princeton University Press, Princeton, NJ.

7 Bierman H. and Hausman W. H. (1970). The credit granting decision. *Management Science.* **16 (8):** B519-B532.

8 Breiman L. (2001). Random Forests. *Machine Learning.* **45 (1):** 5-32.

9 Breiman L., Friedman J., Olshen R. and Stone C. (1984) *Classification and Regression Trees,* Wadsworth, Monterey, CA.

10 Brugha C. M. (2004). Phased multicriteria prefence finding. *European Journal of Operational Research.* **158:** 308-316.

11 Choi T.-M., Li D. and Yan H. (2003). Optimal two-stage ordering policy with bayesian information updating. *Journal of the Operational Research Society.* **54:** 846-859.

12 Crone S. F., Lessmann S. and Stahlbock R. (2004). Empirical comparison and evaluation of classifier performance for data mining in customer relationship management. *International Joint Conference on Neural Networks (IJCNN'04).* Budapest, Hungary. **1:** 443-448.

13 Crone S. F., Lessmann S. and Stahlbock R. (To be published). The impact of preprocessing on data mining: an evaluation of classifier sensitivity in direct marketing. *European Journal of Operational Research.*

14 Crook J. N., Edelman D. B. and Thomas L. C. (2001). Editorial overview. *Journal of the Operational Research Society.* **52:** 972-973.

15 Denardo E. V. (1982) *Dynamic programming: models and applications,* Prentice-Hall, Inc., Englewood Cliffs, N. J. 07632.

16 Desai V. S., Crook J. N. and Overstreet Jr. G. A. (1996). A comparison of neural networks and linear scoring models in the credit union environment. *European Journal of Operational Research.* **95:** 24-37.

17 Ekárt A. and Németh S. Z. (2005). Stability analysis of tree structured decision functions. *European Journal of Operational Research.* **160:** 676-695.

18 Geoffrion A. M. and Krishnan R. (2003). E-business and management science: Mutual Impacts. *Management Science.* **49:** 1275-1286.

19 Gittins J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B.* **41 (2):** 148-177.

20 Gittins J. C. (1989) *Multi-armed bandit allocation indices,* John Wiley & Sons Ltd., Great Britain.

21 Gittins J. C. and Jones D. M. (1974) In *Progress in Statistics*(Eds, J.Gani, K.Sakadi and I.Vinczo) North Holland, Amsterdam, pp. 241-266.

22 Gönül F. and Shi M. Z. (September 1998). Optimal mailing of catalogs: a new methodology using estimable structural dynamic programming models. *Management Science.* **44 (9):** 1249-1262.

23 Haddadi S. and Ouzia H. (2004). Effective algorithm and heuristic for the generalized assignment problem. *European Journal of Operational Research.* **153:** 184 - 190.

24 Harper P. R., Sayyad M. G., de Senna V., Shahani A. K., Yajnik C. S. and K.M.Shelgikar (2003). A systems modelling approach for the prevention and treatment of diabetic retinopathy. *European Journal of Operational Research.* **150 (1):** 81-91.

25 Harper P. R. and Winslett D. J. (2006). Classification trees: A possible method for maternity risk grouping. *European Journal of Operational Research.* **169 (1):** 146-156.

26 Häubl G. and Trifts V. (Winter 2000). Consumer decision making in online shopping environments: the effects of interactive decision aids. *Marketing Science.* **19 (1):** 4-21.

27 Haughton D. and Oulabi S. (1997). Direct marketing modeling with CART and CHAID. *Journal of Direct Marketing.* **11 (4):** 42-52.

28 He J., Shi Y. and Xu W. (January 2005). Classifications of credit cardholder behavior by using multiple criteria non-linear programming. *Lecture Notes in Computer Science.* **3327:** 154-163.

29 Heilman C. M., Kaefer F. and Ramenofsky S. D. (2003). Determining the appropriate amount of data for classifying consumers for direct marketing purposes. *Journal of Interactive Marketing.* **17 (3):** 5-28.

30 Holloway H. A. and White C. C., III (2003). Question selection for multi-attribute decision-aiding. *European Journal of Operational Research.* **148:** 525-533.

31 Jung K. and Thomas L. (2004). A note on Coarse Classifying in Acceptance Scorecards. *Discussion Papers in Management, University of Southampton.* **M04-16 (February 2004):** 14.

32 Knott A., Haynes A. and Neslin S. A. (2002). Next-product-to-buy models for cross-selling applications. *Journal of Interactive Marketing.* **16 (3):** 59-75.

33 Lee T.-S., Chiu C.-C., Chou Y.-C. and Lu C.-J. (February 2006). Mining the customer credit using classification and regression tree and multivariate adaptive regression splines. *Computational Statistics and Data Analysis.* **50 (4):** 1113-1130.

34 Lewis E. M. (1992) *An introduction to credit scoring,* The Athena Press, San Rafael, California.

35 MacQueen J. (1966). A modified dynamic programming method for Markovian decision problems. *Journal of Mathematical Analysis and Applications.* **14:** 38-43.

36 Mehta D. (October 1968). The formulation of credit policy models. *Management Science.* **15 (2):** B30-B50.

37 Meyer R. J. and Shi Y. (May 1995). Sequential choice under ambiguity: intuitive solutions to the armed-bandit problem. *Management Science.* **41 (5):** 817-834.

38 Mitchell T. M. (1997) *Machine Learning,* McGraw-Hill, Boston, Mass.[u.a].

39 Montgomery A. L. (2001). Applying quantitative marketing techniques to the internet. *Interfaces.* **31 (2):** 90-108.

40 Murthi B. P. S. and Sarkar S. (2003). The role of the management sciences in research on personalization. *Management Science.* **49 (10):** 1344-1362.

41 Myers J. H. and Forgy E. W. (September 1963). The development of numerical credit evaluation systems. *Journal of American Statistical Association.* **58:** 799-806.

42 Ohlson J. A. (Spring 1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research.* **18 (1):** 109-131.

43 Papamichail G. P. and Papamichail D. P. (To be published). The *k*-means range algorithm for personalized data clustering in e-commerce. *European Journal of Operational Research.*

44 Peltier J. W. and Schribrowsky J. A. (Fall 1997). The Use of Need-Based Segmentation for Developing Segment -Specific Direct Marketing Strategies. *Journal of Direct Marketing.* **11 (4):** 53-62.

45 Piramuthu S. (2004). Evaluating feature selection methods for learning in data mining applications. *European Journal of Operational Research.* **156:** 483-494.

46 Porteus E. L. (2002) *Foundations of Stochastic Inventory Theory,* Stanford Business Books, Stanford, California.

47 Pöyhönen M., Vrolijk H. and Hämäläinen R. P. (2001). Behavioral and procedural consequences of structural variation in value trees. *European Journal of Operational Research.* **134:** 216-227.

48 Prinzie A. and Van den Poel D. (2006). Investigating purchasing-sequence patterns for financial services using Markov, MTD and MTDg models. *European Journal of Operational Research.* **170 (3):** 710-734.

49 Purang K. (2005) Mainstream lenders tap into the non-standard client base in *Datamonitor Report* London.

50 Puterman M. L. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming,* John Wiley & Sons, New York.

51 Raghu T. S., Kannan P. K., Rao H. R. and Whinston A. B. (2001). Dynamic profiling of consumers for customized offerings over the Internet: a model and analysis. *Decision Support Systems.* **32:** 117-134.

52 Rossi P. E., McCulloch R. E. and Allenby G. M. (1996). The Value of Purchase History Data in Target Marketing. *Marketing Science.* **15:** 321-340.

53 Sörensen K. and Janssens G. K. (2003). Data mining with genetic algorithms on binary trees. *European Journal of Operational Research.* **151:** 253-264.

54 Srinivasan V. and Kim Y. H. (July 1987). Credit granting: a comparative analysis of classification procedures. *The Journal of Finance.* **42 (3):** 665-681.

55 Stump R. L., Athaide G. A. and Joshi A. W. (2002). Managing seller-buyer new product development relationships for customized products: A contingency model based on transaction cost analysis and empirical test. *The Journal of Product Innovation Management.* **19:** 439-454.

56 Thomas L. C. (2000). A survey of credit and behavioral scoring: Forecasting financial risk of lending to customers. *International Journal of Forecasting.* **16:** 149-172.

57 Thomas L. C. (July/August 1994). Application and solution algorithms for dynamic programming. *The Institute of Mathematics and its Applications-Bulletin.* **30 (7/8):** 116-122.

58 Thomas L. C., Edelman D. B. and Crook J. N. (2002) *Credit scoring and its applications,* Society for Industrial and Applied Mathematics (SIAM), Philadelphia.

59 Thomas L. C., Jung K. M., Thomas S. D. and Wu Y. (April 2006). Modelling consumer acceptance probabilities. *Expert Systems and their Applications.* **30 (3)**: 499-506.

60 Thrasher R. P. (Winter 1991). CART: A recent advance in tree-structured list segmentation methodology. *Journal of Direct Marketing.* **5:** 35-47.

61 Van den Poel D. and Buckinx W. (2005). Predicting online-purchasing behaviour. *European Journal of Operational Research.* **166 (2):** 557-575.

62 Van Gestel T., Baesens B., Suykens J. A. K., Van den Poel D., Baestaens D.-E. and Willekens M. (To be published). Bayesian kernel based classification for financial distress detection. *European Journal of Operational Research.*

63 Whittle P. (1998). Restless bandits: activity allocation in a changing world. *Journal of Applied Probability.* **A25:** 287-298.

64 http://www.management.soton.ac.uk/staff/fairisaacs/