Mistakes in the Loan Market: A One-Armed Bandit Model*

Augusto Nieto-Barthaburu

Department of Economics

University of California, San Diego

July 2008

Abstract

This paper analyzes the situation of a loan officer that makes sequential decisions on whether to grant loans or not to applicants. Before making each decision, the loan officer can observe some characteristic of the applicant, and in the case that a loan is granted, he can observe whether it was paid back or not. On the other hand, when a loan is denied the officer cannot observe the applicant's behavior. This selection problem will have an effect on the lender's ability to learn about the population of borrowers. We find that in some cases the lender will be able to make correct decisions in the long run, but in other cases not, i.e. he can make mistakes forever. The crucial factor that determines whether full learning will occur is the nature of the information that the lender observes from the applicant before making each loan. If this information is in the form of a continuous variable we can expect the loan officer to make correct decisions in the long run. On the other hand, if the observed information is discrete in nature, the loan officer may make mistakes, and those mistakes may persist even in the long run.

*I am deeply indebted to Ross Starr for very helpful discussions.

1 Introduction

The objective of this research paper is to recognize and address a problem that arises naturally in those settings in which decisions have to be made under uncertainty, and the result of those decisions affect the ability of the decision maker of observing the outcome of the variable of interest when it is realized. In particular, consider a loan officer who has to decide whether to grant or deny loans to a group of applicants. Assume that those decisions are made sequentially, and the next decision will be made after the expiration of the current loan. Abstracting from differences in borrowers characteristics (to which we will go back below), the loan officer problem is at any time to decide to grant the loan or not, based in his current beliefs about the probability that the loan will be repaid. Now, granting a loan implies two things: First, the officer takes a lottery, that will have a positive payoff if the loan is repaid and a negative payoff if it is defaulted on. Second, by granting a loan the officer will also be able to observe the behavior of the borrower (i.e. if he pays the loan back or not) and can use that information to update his beliefs about the probability that a loan is paid back. Therefore, there is an extra value in granting a loan beyond that of the lottery that it implies, and that is the value of the new information gained by being able to observe the response of the borrower. On the other hand, when a loan is denied the lottery taken has a certain payoff, and no new information is gained about the distribution of applicants.

The problem exposed above is of the form of a general class of problems known in the statistical literature as the "Two-armed Bandit Problem". Two armed bandit problems were introduced first in statistics, following ideas laid down in Thompson (1933). Early treatments of the two-armed bandit problem in statistics can be found in Bradt, Johnson and Karlin (1954) and Bellman (1956). Berry (1972) provides a good summary of the literature. In its general form, the two-armed bandit problem consists of an experiment design in which at each stage the experimenter has to decide to sample from either of two binary distributions, and can observe the result of each trial (i.e. a success or a failure) before deciding from which distribution to sample the next time. The name given to the problem comes from gambling, since this experiment design is analogous to the problem of a gambler that can play in either of two slot machines, both of them with two possible outcomes, success or failure, and with unknown probabilities of success. The gambler pulls one of the arms, observes the outcome, and then decides which arm to pull next.

In economics bandit models have been used for example to model decisions in labor markets and pricing under demand uncertainty. Rothschild (1974) uses a two-armed bandit type specification to model the situation of a store owner deciding which price, between two possible prices, to charge to the next customer to enter the store. As different stores face a different sequence of buyers, the model explains how a price distribution can be observed in a given market for the same good, a deviation of the "law of one price".

Our problem, described in the first paragraph, is a special case of the two-armed bandit problem, since one of the arms has a known outcome (i.e. if a loan is denied the outcome is known). This special case has been named the "One-armed Bandit" problem, and was studied by Bradt, Johnson and Karlin (1954). This specification is simpler than the general problem, and has the appeal that many intuitive properties, that do not in general hold for the two-armed bandit case, are true in this special case.

In this paper we concentrate on the application of the loan officer presented above. It should be noticed though that this problem is identical in form to other interesting problems in economics. In general, the results in this paper carry over to other situations in which the decision making process is such that in the situation that a negative decision has been made, no new information is gained about the underlaying probability distribution of the individual that the decision was made about. Another important application is university admissions. Typically, in university admissions, at least at a first stage, applicants are evaluated using a small number of characteristics, as SAT score and high school GPA for example. The university sets a cutoff in the SAT-GPA space, which is usually interpreted as the cutoff line above which the probability of obtaining a certain minimum GPA or better in college is greater or equal to a given preset target, and rejects all remaining applicants. The problem with this procedure, as stated above, is that then the university can only observe the performance of the accepted individuals, but not of the rejected. This problem is very similar to the one developed here, and most of the insight gained in our problem can be carried over to it.

The problem of correctly identifying "credit worthy" applicants has been previously approached from an econometric point of view in the credit scoring literature. Altman (1968) is usually credited with pioneering the use of statistical methods in credit scoring. Boyes, et. al. (1989), Crook and Banasik (2002), Feelders (2002), Fortowsky and LaCour-Little (2001), Kraft, et. al. (2003) and Ladd (1998) are some examples. A good review of some of those methods can be found in Altman (1981).

In this paper we work in a setting in which the loan officer can observe a random variable X_i for applicant $i = 1, 2, 3, ...$, that distinguishes applicants from one another. In particular, we think of X_i as being some borrowers' characteristic, as for example race, gender, income, household size, occupation, education, etc., that the loan officer can use in order to decide whether to grant a loan or not. In our setting the lender observes X_i for applicant i, decides whether to grant the loan or not, and then observes if the loan was paid back in the case that it was granted.

The objective of the current paper is to investigate the consequences of this decision making setting, in which no information can be obtained from the rejected individuals, on the quality of the decisions that the loan officer will be expected to make. In particular, we can ask the question: can we expect that credit worthy individuals (in a sense that will be made precise later) will always receive loans in the long run, and on the other hand, individuals not credit worthy will be rejected in the long run?

In fact, we find that the answer to the previous question depends on the nature of the variable X_i. Below we consider two cases, when X_i are qualitative variables, such as race or gender, and when X_i are quantitative variables, as for example income or age. The main result of section 3 states that when X_i are qualitative, the loan officer may make mistakes, even in the long run. Those mistakes are such that it is possible that a certain group of the population of applicants with a given value of X_i, although profitable to be lent to, may be rejected from some moment on, and will not be able obtain loans ever again.

The intuition behind that result is as follows: Assume that at a given stage the state of information of the loan officer dictates that an applicant with certain characteristics has to be denied a loan. Then in the next stage an applicant with the same characteristics should also be rejected, since no new information was gained from the rejected individual (i.e., the state of information at the following stage will be identical, and consequently will lead to the same decision). Now bear in mind that individuals with those characteristics may, as a group, have a high enough probability of returning the loans (in the sense that it would be profitable to be lent to). Therefore if mistaken decisions are made about a certain group of applicants, those mistakes will be made subsequently.

The problem in this case arises because a "bad realization" in the sequence of applicants may occur for that group, making the loan officer believe that their probability of paying back is lower than the truth (e.g., unless the probability of payback for a group is equal to 1, there is a positive probability that the first N applicants from that group default). Then the loan officer will form beliefs about these applicants, and at a certain point it may become optimal for him to stop lending to that group. From that stage on no new information will ever be collected, and the group will never receive loans again.

In section 4 we consider the case of X_i being quantitative. In this case there will usually be an order in X_i, as for example with income, so applicants can be ranked by their value of X_i. In general when this is the case, it will be natural to expect that the probability of a loan being paid back, conditional on X_i, will change somehow smoothly with X_i. For example, we would not expect that the probability of pay back is substantially different for an applicant with an income of $25,000 than for an applicant with an income of $24,999. In view of this observation, in section 4 we will assume some smoothness of the probability that a loan is paid conditional on X_i, as X_i changes. Under this assumption, we will show that the kind of mistakes that the loan officer could make in the case of X_i being qualitative will not be expected to be made in this case.

The intuition for this result goes as follows: As was pointed out above, in the case of qualitative X_i, some group that has been excluded from getting loans will never be included again. On the other hand, when X_i is quantitative, even when a group has stopped getting loans, the loan officer may, in the future, return to make loans to it. It is true that no new information is being gathered about those applicants, but in this case information being collected about applicants with other values of X_i can be useful to update beliefs about the first group. The following example shows a situation in which this is the case:

Example 1 Without loss of generality, assume that the probability of pay back is monotonically increasing in X_i. In this case, typically at any stage the decision rule of the loan officer will have the form "grant the loan to applicants with a value of X_i greater or equal than a certain value". For instance, assume that in reality the minimum income that would be required to grant a loan is \$23,000, but because of previous experience at a certain stage the decision rule dictates "grant the loan to applicants with income of \$25,000 or more". This rule will of course prevent the lender to get information about individuals with \$24,999 of income. Now, if that rule is used for some time, eventually the loan officer will learn (under certain regularity conditions) that the probability of a loan being paid back for incomes of \$25,000 is in excess of the minimum probability required to grant a loan. As the loan officer knows that the conditional probability does not change much with X_i, he can then use that information to revise the decision rule, and include some individuals with less than \$25,000 of income.

By correcting the decision rule using information gathered for individuals close to the cutoff the loan officer will then be able to include applicants that were incorrectly excluded before, something that was not possible when X_i was qualitative. The main result of section 4 will be that when X_i is quantitative, under some regularity conditions on the underlying conditional probability function, the loan officer will be expected to make correct decisions in the long run. The sense in that those will be correct decisions is that the cutoff in the variable X_i will be correctly set, and only applicants with a high enough probability of paying the loan back (as to be profitable to be lent to) will be granted loans.

A word has to be said here about the implications of these results. The denial of loans to groups that deserve them can be interpreted as discrimination. The reason is that applicant groups that are credit worthy are being denied loans on the base of belonging to that group. The cause, of course, is that unfortunate previous experience makes the lender do so. So, even when not being moved by any discriminatory motive, a profit maximizing lender may exclude a group of individuals from loans, and since no new information will be gathered about these individuals from then on, the group will be excluded forever. As we pointed out above that can be the case when X_i is qualitative, but not when it is quantitative. Therefore the message of these results is that when lenders are profit maximizers, “discrimination” is much more likely to occur related to a qualitative variable, such as race or gender, than related to a quantitative variable, as for example income or age (quotation marks are used to indicate that this is not discrimination in the strict sense of the word). This suggests that the only cases in which government intervention may be justified is when discrimination is related to variables to the first type. In that case, intervention may take the form of temporarily subsidize loans to a certain group, until enough information can be collected about them to be confident about their probability of returning the loan. After that point, profit maximizing lenders can be expected to make correct decisions, and therefore subsidies can be removed. Bear in mind that here we are not considering the case in which the government wants to favor some particular group, and instead the only reason for intervention is to solve the problem that can arise due to lack of information about that group.

It should also be mentioned that the general problem of inability to observe the outcome of certain decisions can also be tackled from an econometric point of view. In recent years the there has been a growing interest in the economic forecasting literature towards taking into account the intended use of the forecasts (i.e. decision making) at all stages of the forecasting practice. In particular, the observation that traditional forecasting loss functions do not necessarily produce optimal forecasts for a given use is central in this new area of research, as pointed out by Granger and Pesaran (2000a, 2000b), Lieli and Elliott (2004), Pesaran and Skouras (2002) and Pesaran and Timmermann (2004) among others. As some of these papers argue, the correct measure to compare the ex-post performance of different forecasts has to be based on the preferences of the decision maker, and they propose statistics to perform such comparisons or evaluations. What has been overlooked so far in this literature is the fact that when decisions are made, in some cases, those decisions will affect the ability of the forecaster to observe the realizations of some of the variables of interest, and therefore will affect his ability to perform the desired comparisons. This angle of the problem will not be explored in the current paper, but appears as interesting for future research.

The paper will be organized as follows: In section 2 we present the basic framework that will be used to study the problem. In section 3 we consider the case when X_i is a qualitative variable, and we state and prove the results mentioned above. In section 4 we tackle the case of X_i being quantitative, and show how in that case the problems that could arise in section 3 are not present. Finally, section 5 contains the conclusions and directions for further research.

2 Basic Framework

To make our problem concrete, say that there is a population of individuals, and for each of them the loan officer will make a decision that will affect an ex-post payoff. Each individual is characterized by a pair (X_i, Y_i), where X_i is a vector of covariates or characteristics of each individual and Y_i is the outcome from the experiment (i.e. it is a binary random variable that takes the value 1 if the individual will pay the loan and 0 if the loan will be defaulted on). Notice that if the loan is denied, the individual does not have the chance of repaying it or defaulted it on. However, we assume that there exists a random variable Y_i that will determine whether the loan will be repaid if it is granted. If the loan is not granted for individual i, then the loan officer cannot observe the value of the variable Y_i. We assume that (X_i, Y_i) are *iid* for $i = 1, 2, ...$

Next, we assume that for each individual ($i = 1, 2, ...$) the decision maker has a utility (payoff) function

$$U(a_i, Y_i)$$

where a_i is an action to be taken, in our case whether to grant a loan or not. Assume that a_i is binary variable, taking the value 1 if the loan is granted and 0 if it is not. At this point it is important to make clear that we are assuming here that the different decisions do no not affect the outcome, just the ability of the officer to observe it. Formally, we say that $p(x) = \Pr[Y = 1 \mid X = x]$ does not depend on a (notice that we have dropped the index i since we are assuming that (X_i, Y_i) are *iid*). In terms of our example, this assumption of $p(x)$ not depending on a will be true for instance in the case of small consumption loans, in which case the fact that the individual gets the loan will not affect either his ability nor

Table 1: The loan officer's utility function

	$a = 0$	$a = 1$
$Y = 0$	0	z
$Y = 1$	0	w

willingness to repay.

Given the binary character of the variables a_i and Y_i, we can assign values to the utilities for each combination of a_i and Y_i. Notice, however, that when the loan is denied ($a_i = 0$), utility will be the same whether the applicant would have paid the loan back or not. We can tabulate utility values as in Table 1, where we normalized the utility of not granting the loan to zero. We assume that $w > 0 > z$.

In our model, applicants arrive sequentially, and the loan officer can observe the results of the last loan before deciding whether to grant the next one or not. We assume that the officer objective is to maximize the discounted sum of the payoffs (utilities) from the loans. The discount factor will be b, such that $0 < b < 1$.

In the basic framework, we will assume that all applicants have identical characteristics, i.e. the value of the vector X_i is the same for all applicants. This will be our benchmark case. The analysis of this case is interesting since it can be used as a benchmark against which more general cases can be compared.

The problem of the loan officer at any stage is to decide whether to grant the loan or not to the applicant in turn. Granting the loan amounts to taking a lottery that will pay w with probability p and z with probability $(1 - p)$. Denying the loan implies getting 0 for sure. The parameter p is not known to the loan officer. Instead, we assume that he has prior beliefs about it represented by the probability distribution function $F(p)$.

If the problem had only one stage, the officer will want to grant the loan whenever $wp + z(1 - p) \geq 0$. In a multi-stage setting, however, there is an additional value to granting a loan, beyond that of the lottery that action implies, and that is the value of the new information gained about the distribution of p. This will be key in our problem, since the loan officer may sometimes give loans whose (utility) expectation is negative, just because the value of information is lager than the expected loss in that stage.

Before we can characterize the officer decision problem, we need to introduce some notation. Define

$V_{m,n}$: Expected return obtained proceeding optimally, after $m + n$ loans have been granted, of which m have been repaid and n have been defaulted on.

$F_{m,n}$: Updated beliefs about p, starting from initial beliefs F and after $m+n$ loans have been granted, of which m have been repaid and n have been defaulted on.

With these definitions, we can write an expression for $V_{m,n}$ for the cases in which a loan is granted and in which a loan is denied. In case in the current stage a loan is granted, then

$$V_{m,n} = \int_0^1 pdF_{m,n}(p)[w + bV_{m+1,n}] + \int_0^1 (1-p)dF_{m,n}(p)[z + bV_{m,n+1}] \tag{1}$$

The interpretation of this equation is the following: The mean probability of the loan being paid back, conditional on having beliefs $F_{m,n}$, is $\int_0^1 pdF_{m,n}(p)$. In that case, then the officer gets a payoff of w in this stage, and in the next stage (when behaving optimally) the expected future payoff will be $V_{m+1,n}$, since at that time one more loan has been paid back. Similarly the mean probability of a default, conditional on having beliefs $F_{m,n}$, is $\int_0^1 (1-p)dF_{m,n}(p)$. In that case the officer gets a payoff of z in this stage and in the next stage the expected future payoff will be $V_{m,n+1}$.

Similarly to equation (1), we can write an expression for the value of denying the loan at the current stage as follows:

$$V_{m,n} = 0 + bV_{m,n} \tag{2}$$

which implies that $V_{m,n} = 0$.

From equations (1) and (2) we conclude that

$$V_{m,n} = \max\{0, \int_0^1 pdF_{m,n}(p)[w + bV_{m+1,n}] + \int_0^1 (1-p)dF_{m,n}(p)[z + bV_{m,n+1}]\} \tag{3}$$

We can now state and prove the first result of the paper:

Theorem 1 *If a loan officer that behaves optimally denies a loan at any stage, then he will deny loans at every stage after that.*

Next, we present the proof of theorem 1.

Proof: Assume that at a certain stage, after $m+n$ loans have been granted, where m have been repaid and n have been defaulted on, the optimal decision is to deny the current loan, i.e.

$\int_0^1 pdF_{m,n}(p)[w+bV_{m+1,n}] + \int_0^1 (1-p)dF_{m,n}(p)[z+bV_{m,n+1}] < 0$

but contrary to the theorem, after r denials optimal behavior allows to grant a loan. Then the expected payoff of this path is

$0+b0+b^2 0+\ldots+b^{r-1}0+b^r\{\int_0^1 pdF_{m,n}(p)[w+bV_{m+1,n}] + \int_0^1 (1-p)dF_{m,n}(p)[z+bV_{m,n+1}]\}$

But notice that this expression is negative by hypothesis, since $\int_0^1 pdF_{m,n}(p)[w+bV_{m+1,n}] + \int_0^1 (1-p)dF_{m,n}(p)[z+bV_{m,n+1}] < 0$. On the other hand, denying the loan for the r^{th} period yields an expected payoff of 0. Then the proposed behavior is not optimal and a loan must be denied in all future stages. Q.E.D.

The intuition behind the previous result goes as follows: We mentioned above that if the loan is conceded the extra information will be used to update previous beliefs about p, whereas if the loan is denied, no extra information will be collected. It follows that if a loan is ever denied, it should be denied forever (since the state of information was such that the loan was to be denied and such information did not change). Then in some cases a loan will be denied forever if the information gathered so far is bad enough. Note, however, that nothing ensures that this is the correct decision. In fact, we will see below that loans could be denied forever even when the probability of repay p is such that it will be profitable to grant loans to all applicants, i.e. when p is such that $wp+z(1-p)\geq 0$.

The previous result has another interesting implication, one that has been discussed in the literature on credit scoring. In the past has been some interest in analyzing if there is discrimination in the loan markets, in particular against some racial group. One can argue that in a profit driven market discrimination will not occur, and profitable loans will be granted. But as the above theorem shows, even a loan officer that is purely driven by profits may in fact discriminate against some group (i.e. loans will be denied to them even when the group as a whole is credit worthy, in the sense of being profitable to be lent to). The key here is that even when those borrowers are credit worthy, bad experience from the past may make the loan officer decide to deny loans at some stage. and when that happens, no new information will be collected that will make him change his mind in the future.

In section 3 we will extend the basic model to the case in which not all applicants are the same, but instead the loan officer can observe some characteristic of them before making the decision of granting or not the loan.

3 Different Types of Borrowers: The Qualitative Characteristic Case

In this section we analyze the case in which borrowers are not identical, as assumed above, but rather there is a characteristic, that we call X_i for $i = 1, 2, ...$ that distinguish different borrowers. For example, X_i could be income, race, marital state, number of children, etc. In general loans will be evaluated based on a number of characteristics, and not a single one. Here, however, we will assume that X_i is a single random variable. The reason for this assumption is that we want to investigate the difference that X_i being qualitative as opposed to quantitative makes on the quality of the decisions that the loan officer will make. In fact, we will show below that when X_i is qualitative, the loan officer will be exposed to making the same mistakes that we encountered in section 2. In this case, the loan officer can use the knowledge about X_i in making decisions, but we will see that this knowledge will only protect him for making mistakes for some groups of borrowers, and not for all of them. On the contrary, when X_i is quantitative, we will show that the officer will be expected to make correct decisions in the long run.

In this section we will study the case in which borrowers are distinguished by a single characteristic X_i, which is a discrete random variable. We assume that before deciding whether to grant or not a loan to applicant i, the loan officer can observe X_i. We will assume that X_i is binary for all i, taking the values 0 and 1. This assumption is only made for simplicity, and the results obtained will be true when X_i can take a finite number of different values.

We think of X_i as being a qualitative random variable, such as race or sex. In real credit scoring, there are some laws that prohibit lenders from using certain characteristics to make decisions on granting loans. Here we allow the loan officer to use this information. Then our treatment can be viewed in either of two ways: a) X_i is one of the variables that the lender can use to separate borrowers, or b) X_i is one of the variables that cannot be used, in which case we would be thinking of what would be the situation if not such laws would be in place.

From these interpretations the second appears to be the most interesting. This is so in view of controversy about the laws that aim to prevent discrimination in the credit markets. The controversy is about whether the market will work well without regulation (in the sense that individuals who are profitable to be lent to will get loans) or on the contrary if the market solution could involve discrimination, as for example racial discrimination. The argument in favor of the free market is that profit driven lenders will not want to discriminate against any group from which they can make a profit. We will see, however, that this may not be the case. The key here is, as pointed out above, that when bad experience is collected about certain group (that may be "good" borrowers as a group), the lender may, behaving optimally, stop lending to it. And in certain cases this could imply never getting new information about those individuals. We will show below that even in the case in which there are other groups getting loans, the information about these groups may not be enough for the officer to correct his mistakes, the "discriminated" group may not receive loans again.

Although this situation may look as discrimination, it should be pointed out that it arises from pure profit driven behavior. While it is true that "credit worthy" groups of individuals may be rejected, that result does not arise from any group preference (as for example racial preference) from the lender, but from its desire to maximize profits. In this sense this situation does not fit into the definition of discrimination. And in fact it could explain why groups of good borrowers may be excluded from loans in a perfectly profit driven market.

As indicated above, we assume that there is a random variable X_i for each applicant $i = 1, 2, ...,$ that can take the values 0 or 1. Then from the perspective of the loan officer, each applicant is viewed as a pair (X_i, Y_i), where X_i is some borrower characteristic and Y_i is a random variable that will be 0 if the applicant will default on the loan (in case of it being granted) and will be equal to 1 if the loan will be paid back. It is important to note that X_i does not enter the lender's utility function, and the only value of observing it for the officer is to help him learn something about the probability of $Y_i = 1$ for the particular applicant. Throughout the paper we will assume that (X_i, Y_i) are *iid* for $i = 1, 2, ...$

We use the following notation: $p_0 = \Pr[Y = 1 \mid X = 0]$ and $p_1 = \Pr[Y = 1 \mid X = 1]$. The loan officer's beliefs about (p_0, p_1) are represented by the joint distribution function $F(p_0, p_1)$. Next we introduce the following assumptions

(A.3.1) The loan officer's beliefs about (p_0, p_1) are such that p_0 and p_1 are regarded as independent, so we write $F(p_0, p_1) = F_1(p_0)F_2(p_1)$, where F_0 and F_1 are the marginal distributions (beliefs) of p_0 and p_1 respectively.

(A.3.2) Initial beliefs about how X affects the probability of $Y = 1$ (subscripts have been dropped since (X_i, Y_i) are *iid*) are such that

$$\int_0^1 p_1 dF_1(p_1) > \int_0^1 p_2 dF_2(p_2) \tag{4}$$

In words, this assumption amounts to say that the loan officer regards customers with $X = 1$ as more reliable than customers with $X = 0$. This belief may come from some previous experience or prejudice, although we don't need to make this precise here.

Assumption (A.3.2) will derive in the following initial behavior rule:

"*If at the initial period a loan is to be granted to an applicant with* $X_i = 0$*, then a loan will also be granted to an applicant with* $X_i = 1$." Furthermore, we will assume

(A.3.3) The previous decision rule will be maintained forever.

Therefore, at any stage, if the state of information is such that an applicant with $X_i = 0$ is to be granted a loan, we constrain the loan officer to grant loans to individuals with $X_i = 1$. What this constraint does is to act as a protection for individuals with $X_i = 1$. The reason is that no matter how bad the realization they may get, they will keep receiving loans as long as $X_i = 0$ individuals do. The only way they will stop getting loans is that first $X_i = 0$ individuals are denied, and subsequently there is bad experience for $X_i = 1$ borrowers.

Now our goal is to show that even in this case, in which we are imposing an artificial constraint that will protect $X_i = 1$ applicants, there is a positive probability that those individuals will be denied loans in the long run, even when they are profitable to be lent to. The intuition for this result resembles that of the last section. Since if treated separately

$X_i = 0$ individuals will be denied forever with some positive probability, and so do $X_i = 1$ individuals, then there will be a positive probability that both of these events happen. Also, note that given the fact that we are imposing an artificial constraint, and that constraint is acting as a protection to the $X_i = 1$ group, this case will represent a lower bound for the general case, in the sense that if we remove the constraint, the probability of the group $X_i = 1$ to be discriminated against will be higher. Therefore showing that this probability is positive for the constrained case suffices to prove that it is positive for the unconstrained case.

The result in this section is contained in the following theorem:

Theorem 2 *Assume that there are two groups of borrowers, indexed by $X = 0$ and $X = 1$, and that (A.3.1), (A.3.2) and (A.3.3) are true . Then there is a positive probability that the group $X = 1$ will be denied loans forever after a certain point.*

Proof: As we are assuming that there is no correlation between the probabilities p_1 and p_2, then each group can be considered separately. In view of this result, the situation reduces to two different cases the case in Theorem 1. Q.E.D.

What this theorem asserts is that even in the case in which the loan officer can observe applicants characteristics before making decisions, when those characteristics are qualitative this will not be enough to prevent the officer from making mistakes with some positive probability, or in other words, this will not be enough to ensure that the loan officer will make correct decisions in the long run with probability 1. The reason is that when characteristics are qualitative, even when the decision maker may have some information about the ranking of characteristics in terms of the probability of pay back, there is no way of quantifying how big is the difference in that probability. As a consequence, the decision maker has to consider both groups separately. Therefore, the main result of the previous section, that once loans are denied to a group then they will be denied forever, still holds in this case.

We will see in the next section that this will not be the case in the situation in which the borrower's characteristic is a quantitative random variable. The key difference is that in that case, the loan officer, however he decides to set his cutoff in terms of X_i, will be able to get information about the behavior of individuals that are "very close" to the cutoff (in the sense of having a value of X_i close to the cutoff value). That information will allow the officer to revise his beliefs down if necessary, and can possibly make him to start granting loans again to individuals with some value of X_i that were not getting loans at some point. In the next section we will explore the case in which X_i are quantitative random variables, and we will see that in that case the loan officer will not be expected to make mistakes in the long run. In fact, it will be shown that when X_i are quantitative, under some monotonicity assumption about $p(x) = \Pr[Y = 1 \mid X = x]$, the loan officer will (in the long run) set the right cutoff (in the X dimension) with probability 1.

4 The quantitative Characteristic Case

In this section we study the case in which borrowers are distinguished by a single characteristic, a random variable X, but as opposed to the last section, that random variable is continuous, as for example household income. We are interested in studying this case since we expect to obtain fundamentally different results than in the previous section. In particular, it will be shown below that when borrowers characteristics are quantitative we will not expect the loan officer to make the same mistakes as in the qualitative characteristics case in the long run.

The intuition goes as follows: In the qualitative case, there is no obvious way of using information about a group of borrowers to make decisions concerning another group. This is so because characteristics are in general qualitative, and in general do not suggest any way in which the probability of pay back should be related to them. Therefore, when a group of applicants is denied a loan, no new information is collected about them, and as a consequence they are to be denied loans forever.

On the other hand, when the characteristic is a continuous variable, the decision making procedure will typically involve setting a "cutoff" value for the variable X such that applicants with values X_i greater than the cutoff will get the loans and applicants with X_i below the cutoff will be denied. Again no new information is being collected for applicants below the cutoff. However, information about individuals that are granted loans but are "close" to the cutoff, in the sense of having a close value of X_i to the cutoff value, can be used to update the information set of individuals on the other side of the cutoff that are also close to it. Therefore, applicants of a certain value X_i can be considered for loans even after they have been denied in the past. In the long run, it will be shown that this feature will lead to the conclusion that the cutoff is set a the right value, in the sense that only applicants that are profitable to be lent to get the loans or that the applicants that are denied are so because they are not good credit risks.

Again, applicants arrive sequentially, and are indexed by $i = 1, 2, 3, ...$ As before, the loan officer can observe the variable X_i before deciding whether to grant the loan or not to applicant i. In this section $X_i \in \mathbf{R}$ for $i = 1, 2, 3, ...$ We denote by $F(x)$ the (marginal) distribution function of the variable X_i, $i = 1, 2, 3, ...$ Remember that (X_i, Y_i) are *iid* for $i = 1, 2, 3, ...$

The lender has initial beliefs $q_0(x)$ about the conditional probability $p(x) = \Pr[Y = 1 \mid X = x]$, for $x \in [a, b]$. Although $p(x)$ is unknown, we will assume that it is monotonically increasing in x for all $x \in [a, b]$. We will assume that $p(x)$ is a smooth function of x. In particular, we will assume that it is differentiable for all x in the interior of the support of X:

(A.4.1) The support of X is the interval $[a, b] \subset \mathbf{R}$.

(A.4.2) $p'(x)$ exists for all $x \in (a, b)$.

(A.4.3) $p'(x) > 0$ for all $x \in (a, b)$.

We can now analyze the behavior of the loan officer. Because of our monotonicity assumption, we know (and so does the loan officer) that $p(b) > p(a)$. Therefore it must be the case that $q_0(b) > q_0(a)$. Assume that the lowest probability of payback required for a loan to be granted is k. To make the case interesting, we will assume:

(A.4.4) $q_0(b) > k$

i.e. according to the officer's initial beliefs, there are some values of X for which applicants are good credit risks. This assumption rules out the uninteresting case in which initial beliefs dictate that no applicant, regardless of it's X value, is credit worthy. In that case, no loan will be ever granted. We assume further that:

(A.4.5) $p(b) > k > p(a)$

or in other words, according to the true conditional probability, there are some applicants that are creditworthy and some that are not. Then there is a value $c \in [a, b]$ such that $p(c) = k$. The value c is the value at which the cutoff would be set if the true conditional probability function $p(.)$ would be known.

In addition to the assumptions on the true conditional probability $p(x)$ stated above, we add the following Lipswich condition:

(A.4.6) There exist $0 < g$, such that $0 < p'(x) \leq g$ for all $x \in [a, b]$.

This condition ensures that the probability of pay back does not change much with X. Remember that we are assuming that $p(x)$ is monotonically increasing in x, as it would be expected for example if the random variable X represents income. In light of this assumption, this is an appealing condition, since we wouldn't expect the conditional probability of payback to take big jumps as we change the borrower's characteristic X by a small amount. For example, we wouldn't expect that applicants with $25,000 of income have a substantially higher probability of pay back than those with $24,999 of income.

The following condition will also be useful in the proof of the main result:

(A.4.7) $\exists \varepsilon > 0$ such that loans will be granted to all applicants with $X_i \in [b - \varepsilon, b]$ for all $i = 1, 2, 3, ...$

This condition will ensure that there is a small interval in the support of the characteristic for which applicants with those values of X will always receive loans. This assumption is not very restrictive either. Remember that we were considering the case in which $p(b) > k$. In that situation, due to the continuity of $p(x)$, we know that there exist $\varepsilon > 0$ such that $p(x - \varepsilon) > k$. Furthermore, in view of the condition that $p'(x) \geq h > 0$, an ε with that property can be easily computed. Therefore the condition amounts to say that loans will be always granted to a group of applicants that are credit worthy.

We will denote by $q_i(x)$ the updated beliefs of the loan officer about the probability $p(x)$, after applicant i has been evaluated, and (if the loan was granted) his behavior has been observed. In view of assumption (A.4.3), for each $i = 1, 2, ...$ the function $q_i(.)$ defines a cutoff value, such that loans will be granted for all x such that $q_i(x) \geq k$. We call that cutoff value t_i, i.e. $q_i(t_i) = k$. In other words, t_i is the cutoff in the income dimension set by the lender before applicant $i+1$ is evaluated. Then the decision rule for the lender will be

"grant a loan to applicant $i+1$ if his income is larger than t_i, and deny him the loan otherwise".

To this point nothing has been said about the updating rule that the loan officer uses to correct his beliefs in view of new information. The typical assumption in the literature is that beliefs are updated using Bayes rule. Here we refrain from assuming a particular updating rule, since different rules may be preferred by different lenders. Instead, we prefer to impose some minimal condition that any sensible updating rule should have. In particular, what the next condition states is that the updating rule used by the loan officer will be such that his beliefs about the conditional probability of payback for a given interval in the support of X will converge (in the sense of convergence almost surely) to the true conditional probability for that interval when the number of observations in the interval grows to infinity. We introduce the following condition on the updating rule used by the loan officer:

(A.4.8) The updating rule used is such that $q_i(x) \rightarrow_{a.s.} p(x)$ as $i \rightarrow \infty$ for all x such that $t_i < x$ infinitely often.

The intuitive argument behind this condition is of course the strong law of large numbers. We know that the variables X_i are *iid* for $i = 1, 2, 3, ...$ Then when we say a "sensible updating rule", the meaning of that is that a law of large numbers applies for all values of x for which enough information is collected, in the sense that the cutoff falls to the left of x for infinitely many values of i. When this is the case, that particular x will be in the interval that will get loans for infinitely many periods.

Our objective here is to show that in this case the loan officer will not make the kind of mistakes that he could make in the qualitative characteristics case. The reason is that since he can get a very good estimate of the conditional probability $p(x)$ in the interval $[b-\varepsilon, b]$, he can then use that knowledge and extend it to other values in the support of X. The key for doing that is the knowledge about the true conditional probability that the officer has, namely that contained in condition 3. We will prove that when the lender has that knowledge, then in the long run he will make correct decisions, in the sense that applicants will be granted loans if and only if $p(x) \geq k$.

The next theorem contains the main result of the section:

Theorem 3 *Assume (A.4.1)-(A.4.8). Let* $t_i \in [a, b]$ *such that* $q_i(t_i) = k$*. Then* $t_i \to c$ *as* $i \to \infty$.

What the theorem claims is that in the long run the loan officer will use the right cutoff, in the sense that it will grant loans to credit worthy applicants and it will deny loans to bad credit risks.

Proof: We will show that $\limsup t_i = \liminf t_i = c$. We first show that $\limsup t_i = c$. Assume this is not the case. Then either $\limsup t_i > c$ or $\limsup t_i < c$. We want to show that in either of these cases the loan officer would be making suboptimal decisions.

Suppose that $\limsup t_i = t < c$.We know that, by monotonicity of $p(x)$, $p(t) < p(c)$. Now pick $\delta > 0$ such that $t + \delta < c$. Since $\limsup t_i < t + \delta$, almost all elements of the sequence $\{t_i\}$ fall in the interval $[a, t + \delta]$. Then there is N such that $t_i \leq t + \delta$ for all $i > N$. But that means that loans will be granted to applicants in the interval $(t + \delta, c)$ for all $i > N$, and therefore $q_i(x) \to_{a.s.} p(x)$ for $x \in (t + \delta, c)$. Now, as for all such x it is the case that $p(x) < p(c) = k$, this will imply that the loan officer is granting loans in the long run to applicants that he learned not to be credit worthy, in the sense of having a lower probability of returning the loan than the minimum probability required to grant a loan. Therefore the officer must be behaving suboptimally.

On the other hand, suppose that $\limsup t_i = t > c$. Again, by monotonicity of $p(x)$, $p(t) > p(c)$. For any $\delta > 0$ such that $t - \delta > c$, $t_i < t + \delta$ for almost all i. Also we know that $t_i > t - \delta$ infinitely often. Then by condition 5 above, since this is true for any δ, $q_i(x) \to_{a.s.} p(x)$ for $x \in (t, b]$, and by continuity of $p(x)$, $q_i(x) \to_{a.s.} p(x)$ for $x \in [t, b]$. Now by condition 3, we can pick δ such that $p(t - \delta) > p(c)$. But then the fact that $t_i > t - \delta$ infinitely often implies that the loan officer will set the cutoff infinitely often above a point for which he will know that the probability of pay back is larger than the minimum required. Therefore he is making suboptimal decisions. Then we have shown that $\limsup t_i = c$.

A symmetric argument shows that $\liminf t_i = c$. Then since $\limsup t_i = \liminf t_i = c$ we conclude that $\lim t_i = c$ and the proof is complete. Q.E.D.

The result contained in theorem 3 is quite surprising. Its interpretation is the following: When the borrowers' characteristic that the loan officer uses is quantitative, then he may make mistakes in the short run, but in the long run all credit worthy individuals (and only those) receive loans. The key behind the result is the ordering inherent to the variable X, that can be translated to the conditional probability of payback. Then the loan officer is able to learn from individuals that are not getting loans, since they are related by X to some individuals that are being approved. That ability to learn from those who are rejected is the main difference with the qualitative case, and is the main reason behind the sharp difference in the results for both cases.

5 Conclusion

The paper considered a decision problem in which a loan officer makes sequential decisions on whether to grant loans or not to individuals who apply for them, based on the observation of some borrowers' characteristics that can be useful to distinguish among different borrowers in terms of their probability of paying the loan back.

It was shown that the kind of information contained in the observable characteristic X_i has an important effect on the quality of decisions that the loan officer will be expected to make in the long run. In particular, we distinguish between two cases, when X_i is a qualitative variable and when X_i is a quantitative variable. In the first case, the nature of X_i does not allow the loan officer from using information from one group of applicants to update his beliefs about another. Given that, if it is the case that a certain point in time the loan officer decides to deny loans to a certain group, that group will be denied loans from that moment on. The reason is because they are denied loans, no new information is being collected about those applicant. Therefore, if the state of information was such that they had to be denied loans at a certain point, it will stay that way from there on, and those individuals will no longer receive loans.

The important point in this case is that this situation of a group being excluded from the market from some moment on may happen even when those individuals are "credit worthy", in the sense of having a high enough probability of paying the loan back as to be profitable to be lent to. The reason why such situation can arise is because the initial realization of the payback sequence for those individuals may be so discouraging that the loan officer may decide to stop “buying” information about those individuals, and therefore will never learn that they are in fact credit worthy.

On the other hand, it was shown that when X_i is a quantitative variable, those problems are not expected to occur. When this is the case, typically we can expect some “smoothness” in the function $p(x)$, the conditional probability of a loan being paid back when $X_i = x$. By smoothness we mean that we don't expect $p(x)$ to change much do to a small change in x. When this is the case, it was shown that in the long run the loan officer will be expected to make correct decisions, in the sense that it will grant loans to individuals for whom $p(x) > k$, where k is the minimum probability of payback that will make granting a loan profitable.

The reason why the problems encountered in the qualitative characteristic case are not present here is because in this case the loan officer, even when not gathering information for a certain group of individuals, will be gathering information for other individuals that can be related to them by proximity in the x dimension. Therefore the lender will be able to correct mistakes made in the past, and possibly include again among those who receive loans some individuals that where rejected in the past.

The results in the paper have strong policy implications. In fact, the message is that if we assume that lenders are profit maximizers, then we would not expect discrimination to occur related to quantitative variables. Therefore, the argument that low income people are being discriminated would not be supported by our results. On the other hand, discrimination could occur related to a qualitative variable, such as race or gender. In that case, intervention may be justified in the form of subsidies to loans for a certain group. Nevertheless, our findings would support such subsidies only for a period of time, until enough information can be obtained about the payback probability for the group in question. Beyond that, profit maximizing lenders behaving optimally should be able to make correct decisions and no discrimination is expected to happen.

It was also mentioned that the problem dealt with in this paper is one application of a general type of problems, where the decision maker cannot observe the realization of certain variable of interest for certain decisions. Another important application is university admissions, since rejected applicants cannot be evaluated and therefore their performance cannot be observed.

In section 3 we assumed that there is no a priori correlation between the probabilities of payback of the groups of applicants. An extension to this work would be to relax that assumption. In that case, our intuition tell us that we could expect that the loan officer will be less exposed to make mistakes in the long run. However, it appears that those results will depend on the knowledge that the loan officer has about the correlation between the probabilities for different groups. This constitutes an interesting line for future research.

It was mentioned above that this problem can be approached for an econometric point of view as well. In particular, there has been interest in recent years in developing forecasting methods that take into account the preferences of the users of the forecasts. Also, there has been interest in developing methods to evaluate and compare forecasts from the point of view of the forecast user, taking as a measure of forecast efficacy the quality of decisions it leads to. What has been overlooked so far in this literature is the fact that when decisions are made, in some cases, those decisions will affect the ability of the forecaster to observe the realizations of some of the variables of interest, and therefore will affect his ability to perform the desired evaluations or comparisons. This angle of the problem appears as interesting for future research.

References

[1] Altman, E. I. (1968): “Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy” Journal of Finance 23: 589-609.

[2] Altman, E. I., R. B. Avery R. A. Eisenbeis and J. F. Sinkey (1981): “Application of Classification Techniques in Business Banking and Finance” JAI Press: Greenwich, CT.

[3] Bellman, R. (1956): “A Problem in the Sequential Design of Experiments” Sankhya 16: 221-229.

[4] Berry, D. A. (1972): “A Bernoulli Two-Armed Bandit” Annals of Mathematical Statistics 43: 871-897.

[5] Boyes W., D. Hoffman and S. Low (1989): “An Econometric Analysis of the Bank Credit Scoring Problem” Journal of Econometrics 40: 3-14.

[6] Bradt, R. N., S. M. Johnson and S. Karlin (1954): “On Sequential Designs for Maximizing the Sum of n Observations” Annals of Mathematical Statistics 27: 1060-1070.

[7] Crook, J. and J. Banasik (2002): “Does Reject Inference Really Improve the Performance of Application Scoring Models?” Working paper. Credit Research Centre, The School of Management, University of Edinburgh.

[8] Feelders, A. J. (2002): “An Overview of Model Based Reject Inference for Credit Scoring” Working paper. Institute for Information and Computing Sciences. Utrecht University.

[9] Fortowsky, E. and LaCour-Little, M. (2001): “Credit Scoring and Disparate Impact. Working paper” Wells Fargo Home Mortgage. Available at fic.wharton.upenn.edu/fic/lacourpaper.pdf

[10] Granger, C. W. J. and M. H. Pesaran (2000a): “Economic and Statistical Measures of Forecast Accuracy” Journal of Forecasting 19: 537:560.

[11] Granger, C. W. J. and M. H. Pesaran (2000b): "A Decision Theoretic Approach to Forecast Evaluation" Working paper. Department of Economics, University of California, San Diego.

[12] Kraft, H., G. Kroisandt and M. Muller (2003): "Assessing the Discriminatory Power of Credit Scores Under Censoring" Working paper. Fraunhofer Institut fur Techno-und Wirtschaftsmatemathik, Kaiserslautern, Germany.

[13] Ladd, H. (1998): "Evidence on Discrimination in Mortgage Lending" Journal of Economic Perspectives 12: 41-62.

[14] Lieli, R. and G. Elliott (2003): "A Flexible Framework for Predicting Binary Variables" Working paper. Department of Economics, University of California, San Diego.

[15] Pesaran, M. H. and S. Skouras (2002): "Decision Based Methods for Forecast Evaluation" In M. P. Clements and D. F. Hendry (Eds.) A Companion to Economic Forecasting. Oxford: Basil Blackwell, 241-267.

[16] Pesaran, M. H. and A. Timmerman (2004): "Real Time Econometrics" CESifo Working Paper No. 1169.

[17] Rothschild, M. (1974): "A Two- Armed Bandit Theory of Market Pricing" Journal of Economic Theory 9: 185-202.

[18] Thompson, W. R. (1933): "On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples" Biometrika 25: 285-294.