INGENIERIA INDUSTRIAL
UNIVERSIDAD DE CHILE

CENTRO DE FINANZAS

Granting and Managing Loans for Micro-Entrepreneurs: New Developments and Practical Experiences

Cristián Bravo,
Sebastián Maldonado
y
Richard Weber.

El Centro de Finanzas cuenta con el significativo apoyo del Banco de Crédito e Inversiones BCI

fcfm

Granting and Managing Loans for Micro-Entrepreneurs: New Developments and Practical Experiences*

Cristián Bravo†✠, Sebastián Maldonado‡, and Richard Weber†

✠Finance Center, Department of Industrial Engineering, Universidad de Chile.
†Department of Industrial Engineering, Universidad de Chile.
‡Universidad de Los Andes, San Carlos de Apoquindo 2200, Las Condes, Santiago, Chile.
cbravo@dii.uchile.cl,smaldonado@uandes.cl,rweber@dii.uchile.cl

Final version: 9 September, 2012.

Abstract

We present a methodology to grant and follow-up credits for micro-entrepreneurs. This segment of grantees is very relevant for many economies, especially in developing countries, but shows a behaviour different to that of classical consumers where established credit scoring systems exist. Parts of our methodology follow a proven procedure we have applied successfully in several credit scoring projects. Other parts, such as cut-off point construction and model follow-up, however, had to be developed and constitute original contributions of the present paper. The results from two credit scoring projects we developed in Chile, one for a private bank and one for a governmental credit granting institution, provide interesting insights into micro-entrepreneurs' repayment behaviour which could also be interesting for the respective segment in countries with similar characteristics.

1 Introduction

Credit scoring corresponds to the use of statistical models to transform relevant data into numerical measures that guide credit decisions (Anderson, 2007), and its main objective is to estimate the probability of default, i.e. the event of a customer not paying back the loan in a given time period. Recent developments in credit scoring are oriented, for example, in analysing imbalanced credit datasets (Brown & Mues, 2012), in survival analysis (Bellotti & Crook, 2008; Tong et al., 2012), in correct ways to validate credit scoring models and make them comprehensible (Castermans et al., 2010), and in adapting new models for credit scoring use (Setiono et al., 2009).

This work focuses on micro-entrepreneurs, a segment different from independent persons or large companies in terms of size, income, and organizational structure. Although several efforts have been undertaken to gain knowledge about the default risk of small and medium-sized enterprises – see, e.g., Kim & Sohn (2010), or Van Gool et al. (2011) –, only few studies are available for micro-entrepreneurs (Schreiner, 2000).

The aim of this paper is two-fold: First, it provides experiences and insights we gained from several credit scoring projects – a private bank and a state-owned organization – where we had to adapt existing methodologies for credit granting. Secondly, these projects required the development of new techniques for credit scoring, namely cut-off point construction and model follow-up, which will be described subsequently.

This paper is organized as follows. The next section characterizes Chilean micro-entrepreneurs. Subsequently, we present the methodology used to construct credit scoring models putting special emphasis on stages where problems arose or special knowledge was revealed. The following section contains the results

*The following is a post-acceptance, pre-copyeditting version of our work to appear at the European Journal of Operational Research, available at www.sciencedirect.com. Please cite this paper as follows: Bravo, C., Maldonado, S., and Weber, R. Granting and Managing Loans for Micro-Entrepreneurs: New Developments and Practical Experiences. European Journal of Operational Research. DOI: http://dx.doi.org/10.1016/j.ejor.2012.10.040

we obtained applying the proposed methodology to the state-owned organization. Finally, conclusions are drawn from our work in Section 6.

2 Chilean Micro-entrepreneurs

In most countries, micro-entrepreneurs are an important element of the economy, accounting for the creation of new business opportunities and employment. Chilean micro-entrepreneurs are defined as very small firms, with up to ten employees and an annual income of no more than EUR 97,000. They offer 21% of the jobs in the country (División Empresas de Menor Tamaño, 2009), and represent a large portion of Chilean companies), but they generate only a small portion of annual sales, with micro-entrepreneurs representing 13.9% of total annual sales of all companies in the country (Instituto Nacional de Estadísticas, 2002).

Micro-entrepreneurs, usually receiving support only from governments and not-for-profit initiatives, have become attractive customers for banks and loan-granting institutions in recent years, but the risk measures that accompany them have to follow suit. In particular, micro-entrepreneurs have certain characteristics that must be taken into account when developing a credit scoring system. For example, they are usually on a tight budget, regardless of their revenues. So it is not a question of how much money they make, which takes away a natural candidate for a discriminatory variable in credit scoring models. Considering the special characteristics that micro-entrepreneurs possess, there is still little knowledge on the variables that may determine their risk as loan borrowers. Additionally, most micro-entrepreneurs have not had any access to financial instruments previously, so the common scorecards that are available from major distributors are concentrated on higher risk segments (due to lack of information), and makes the segment be considered as "high risk" without any deeper consideration (Ministerio de Economía, 2012).

The question may arise of whether funding is even needed for this segment. A recent government study (Ministerio de Economía, 2012) found that 81 percent of micro-entrepreneurs fund their start-ups using their personal savings or family loans. Only four percent of the start-ups were funded using banks or established financial institutions, reporting that the access to credit is given only by supermarkets and some retailers (in the form of personal credit cards). The main conclusion we draw from the information presented is that loans are necessary for this segment, but the perceived risk of the micro-entrepreneur is too high for traditional financial institutions to provide coverage. We believe that existing credit scoring methodologies must be adapted in order to correctly reflect the reality of micro-entrepreneurs and, in that way, create the conditions for an equal-opportunity and profitable environment both for micro-entrepreneurs and for the institutions that grant them loans.

3 Developing a Credit Scoring System

We applied the KDD process (Fayyad et al., 1996) to develop both models; one for the private bank and the other for the state-owned organization. In the following subsection we describe the process of data acquisition and consolidation, followed by the process of data cleansing. Finally we present the process for estimating the probability of default for a solicited loan.

3.1 Data Set Consolidation

The first steps were to identify the relevant databases in which the information was scattered, extract the respective variables, and load them into a repository especially created for the construction of our models.

In order to develop an effective credit scoring system, a homogeneous and representative sample of the population for both classes (defaulters and non-defaulter) is needed. Based on the relevant literature, e.g. Thomas et al. (2002), and our experience from other similar credit scoring projects, we first segmented customers and the requested loans by differentiating between new customers in one group and renewing or current customers in the other.

3.2 Data Cleansing and Variable Selection

The next step consisted of data preparation and variable (feature) selection. The procedure applied had two goals: to select features in a cascade-like approach thus minimizing the risk of eliminating potentially useful ones, and to maximize the knowledge extracted from the respective dataset. The description of the procedure follows:

1. Concentration of feature values and analysis of missing values: In order to quickly discard useless variables, those concentrated in a single value in more than 99% of the cases, and those with more than 30% of missing values were eliminated. The rationale of the second criterion is to reduce the number of discarded cases or imputations realized in order to construct the final model.

2. Univariate analysis: The remaining variables were tested to find distribution equality across groups, using the objective variables (defaulter / non-defaulter) as the splitting criterion. In particular Kolmogorov-Smirnov (K-S) and χ^2-tests were applied to continuous and discrete variables, respectively.

3. Final data preparation: The selected variables accounted for 20% of the original ones. The resulting dataset had very few null cases (less than 1%), which were eliminated, and resulted in a robust subset of the variables which were then used as input for a logistic regression model.

3.3 Estimating the Probability of Default

The final decision regarding the acceptance of the loan application can be made by comparing the calculated probability of default, estimated using a suitable model and the variables obtained from previous steps, with a suitable pre-defined threshold (Hand & Henley, 1997). Several studies in credit scoring have focused on comparing the classification performance of different techniques (see, for example, Baesens et al. (2003), Finlay (2011)). Their main conclusion is that the traditional credit scoring method logistic regression reaches performance levels comparable with more sophisticated data mining approaches for modelling credit risk, being therefore the main technique used for credit scorecard construction (Thomas et al., 2002).

The objective of logistic regression is to construct a function which determines the probability of default for a given client. It considers V different regressors in a vector $\mathbf{x}_i \in \mathbb{R}^V, i = 1, \ldots, N$, and an observed binary variable $y_i = 1$ if borrower i defaults, and $y_i = 0$ else. Considering the dependent variable as latent, the probability of default is:

$$p(y_i = 1|\mathbf{x}_i) = \frac{1}{1 + e^{-\left(\beta_0 + \sum_{j=1}^{V} \beta_j x_{ij}\right)}}, \tag{1}$$

where β_0 is the intercept, and β_j is the regression coefficient associated to variable j. Since these parameters are unknown, they have to be estimated using, e.g., a maximum likelihood algorithm, which results in unbiased, asymptotically normally distributed estimators $\hat{\beta}_j$ (Hosmer & Lemeshow, 2000). The expression in the exponent is a measure of the total contribution of the independent variables used in the model, and is known as the logit (Greene, 1993).

The process of constructing the model uses a wrapper for variable selection, consisting of a greedy search of variables that accounts for maximum discriminatory capacity. There are two common approaches (Hosmer & Lemeshow, 2000): *forward selection* and a *backward elimination*. Both methods are based on measuring how relevant each variable is, as measured by a χ^2-test.

4 New Developments

The final step when constructing a credit scoring system is deciding the cut-off to be used in order to transform the probability obtained from the logistic regression into a binary decision. However, when profits are not well-defined, this is not a trivial task as we found, e.g., in our project with the state-owned, non-profit institution. We developed a generic methodology to determine such cut-off points for the case of non-profit organizations. Additionally, once the credit scoring model is being used, different kinds of changes might occur diminishing predictive capability. We developed procedures to follow such shifts in the population and update the respective model parameters accordingly. These approaches are presented in the subsequent subsection.

4.1 Methodology for Cut-Off Point Construction

The methodology to determine a cut-off point starts by determining the cost of accepting a bad applicant using the expected loss for a given loan. The second part is estimating the cost of rejecting a good applicant, point especially interesting for any governmental institution, usually more interested in public wealth than in monetary profits. Finally, both numbers are considered and combined to calculate an optimal cut-off point to be applied to the result of the logistic regression model.

4.1.1 Cost of accepting a bad applicant

Our cut-off point methodology starts by segmenting the range of the possible values the estimated default probability can take. Considering the database of loans that were used to validate the model, it is necessary to obtain the probabilities of default for each customer. Then these values are ordered and segmented in ranges depending on the size of the database. In our experience, intervals of 0.05 turned out to provide the best results, generating 20 segments of the interval [0.1].

The direct cost for the organization when a grantee does not repay his or her loan has to be calculated by taking into consideration all the resources owed. Additionally, the value of the collateral must be discounted, that is, the real estate or different properties that the grantee declares as security for the loan. In general, the loss per loan (given that default occurred) is the product of the following quantities(Ozdemir & Miu, 2009):

- EAD: Exposure at default. Is the amount that the grantor is owed when default occurs, including all standing instalments and any owed interest. In the case of loans with guarantors, the value of the loss and the exposure is different (Superintendencia de Bancos e Instituciones Financieras, 2008), but for this particular case (no guarantors considered) they are assumed to be the same.

- LGD: Loss Given Default. Proportion of the exposed capital (EAD) that is actually lost given the event of default. This value considers the expected proportion of the loan that will not be paid by the customer after default occurs including the recovery after prosecution or collection, and the recovery given by the collateral.

The cost of defaulting for each customer follows $LOSS = LGD \cdot EAD$. In order to determine the cut-off point it is necessary to know the accumulated cost for all customers. The final amount corresponds to the set of defaulters whose default probabilities are below the cut-off point p:

$$C^p_{\text{Loss}} = \frac{\sum_{i \in D(p)} LOSS_i}{|D(p)|} \qquad p \in \{0.05, 0.1, ..., 0.95, 1\}, \tag{2}$$

where $D(p)$ is the set of defaulters with estimated probability of default less than p, $|D(p)|$ is the number of customers that belong to set $D(p)$, and $LOSS_i$ is the observed loss for customer i.

The cost is divided by $|D(p)|$ so an average cost per loan in that cut-off is determined. This way, when applying the model, corrections can be introduced considering the actual number of loans that are observed – a new $|D(p)|$ – tying the cut-off to market conditions.

4.1.2 Cost of Rejecting a Good Applicant

The case of a private loan-granting company rejecting a potentially good customer leads to an opportunity cost equivalent to the gain or utility the loan would have generated for the financial entity. This cost has an associated market share loss: if the financial policy has been too restrictive and many potentially good borrowers have been rejected, the company exposes itself to a commercial risk by reaching fewer customers than would have been possible.

In the case of a state-owned institution, however, the cost of rejected loans has to consider the profit a rejected loan could generate for the organization (opportunity cost) plus the benefit for society that is not going to occur since the credit is not granted. In general, this benefit is not relevant for private organizations when determining their cut-off points. Furthermore, it could be argued that the profits that any institution (public or private) perceives is at the expense of the borrower, so the net social benefit of profits is zero, therefore should not be considered.

In order to estimate this lost benefit for society, we compared the average income of applicants who **did** receive loans from the institution with average income of entrepreneurs who **did not** receive any loan.

Table 1: Comparison of Annual Income for Customers and Control Group.

	With Loan	Without Loan	Diff./Avg.	Percentage
Avg. annual income	4,869 EUR	3,582 EUR	1,287 EUR	26.44%
Est. loss of income (if loan not granted)	1,247 EUR		1,314 EUR	26.97%
Est. additional income (if loan granted)		1,377 EUR		

A previous study provided by the government institution analysed a sample of 1,010 entrepreneurs who received a loan and a sample of 500 workers who did not receive any loan (control group).

The survey was conducted in the year 2005, considering granted loans between 2000 and 2003 for the study group. These granted loans were repaid before the survey was taken. The survey was controlled by region and activity by using cluster sampling, but also another variables were studied in order to validate the sampling process, such as age, sex, size of the family group, ethnic group, and educational level. No significant differences were found for these variables.

The average annual income obtained by the customers (study group) was 4,869 EUR, which is 1,287 EUR higher than the average of the control group (3,582 EUR). Additionally, the customers who received a loan were asked to estimate their expected loss in terms of annual income in case they would not have received credit. Those who did not receive a loan were asked to estimate how much additional income they would have perceived. The averages of both estimations are similar to the difference of the average annual income for both groups (1,314 EUR and 1,287 EUR respectively), making the estimation robust. Table 1 summarizes the results of this analysis.

According to this information, the average income of applicants with access to a loan is 26.44% higher than the ones without this form of financial help. Given that the vast majority of these entrepreneurs without access to loans do not get credit from private financial institutions and tend to reduce their production level, we consider this percentage to be a valid estimate for the opportunity cost of not receiving a loan.

We assume the opportunity cost as a percentage of the amount of the loan for several reasons: First, the annual income and the amount of the granted loan are highly correlated for micro-entrepreneurs, as we could corroborate for a similar universe of the private bank's customers. This insight is relevant since the variable "income" is not properly available for all applicants. Secondly, we want to provide a comparison between both costs (monetary loss, which depends on the amount of the loan, and opportunity cost) in terms of the same variable. Finally, it is more intuitive to estimate the profit a good payer may generate (and therefore a social benefit) as a proportion of the money that he/she receives (the amount of the loan), instead of the income he/she already has.

The cost of rejecting a good applicant per cut-off point p is:

$$C_{RG}^{p} = 0.2644 L_p,\ p \in \{0.05, 0.1, ..., 0.95, 1\}, \tag{3}$$

where L_p is the average amount of the loans granted for a given cut-off p. It is important to notice that for this particular application no interest rate is considered, and only the effect of the inflation is charged for repayment. In a general case, the interest should be subtracted from the profit generated by the loan.

4.1.3 Cut-Off Point Construction

Given the results of the previous subsections, i.e. knowing the values of C_{Loss}^{p} and C_{RG}^{p}, it is possible to construct the final cut-off point. The optimal value will be simply the one that minimizes the total cost in the test dataset, given by:

$$P_{min} = \text{argmin}_p (C_{\text{Loss}}^{p} \cdot |D(p)| + C_{RG}^{p} \cdot (|G(1)| - |G(p)|)) \tag{4}$$

Similarly to equation 2, $G(p)$ represents the set of customers with probability of default less than p, so $|G(1)|$ represents the total number of customers. For equation 4 we assume a Benthamite welfare function, commonly used for social economics without assuming an uneven distribution of goods, in which an Euro of foregone benefit to society has the same weight as an Euro of loss from default for the financial institution, since the Euro foregone should have been used for a new loan or for improvement of the welfare of a borrower, the damage is indistinct (Arrow & Debreu, 2002).

P_{min} is the cut-off point that minimizes the total cost of misclassification. It should be noted that this expression is valid only if the value of the objective variable used to estimate the logistic regression model is 1 in case the loan defaults, as is recommended when constructing scorecards (Anderson, 2007).

4.2 Model Follow-Up

As in many other projects, the users of our credit scoring systems asked for a way to update the respective models. This practical need has also been recognized in the literature. According to Thomas et al. (2002), a statistical model has an average lifetime of two years before it begins to lose predictive capacity, which directly translates into losses due to a higher default rate as well as higher provisions. A way to prevent the described loss in predictive capacity would therefore be attractive for any credit-granting company, making model follow-up a challenge and an interesting research opportunity.

The need for follow-up in credit scoring also has a very strong regulation component: credit risk models must be first approved by the national supervisor before first use, and this procedure can be time consuming, expensive (since resources must be allocated), and overall stressing the operations of the credit risk area. These facts block simply updating the model every six months, and incentives the institution to wait as much as possible to change the model, usually when losses are already being perceived. Correct follow-up should help in avoiding these losses from taking place. Additionally, the developing a credit scoring model is not just related to estimating the coefficients, since implementing a new model stresses much more than just the risk analysis department in the company. The costs saved in training, systems implementation time and resources, and other overhead costs that arise when a new model is implemented encourage the use of follow-up techniques over implementing a new model from scratch.

Similar to model construction, model follow-up must be easily understandable by its users and easily applicable. Additionally, it must provide an accurate picture of the population shifts. We applied the following three approaches for model follow-up:

1. Variables' Discriminatory Capacity: Over time, variables can lose their discriminatory capacity. This was measured independently of the respective model, using the same procedures performed for variable selection. We applied K-S or χ^2-tests using the predicted result as the splitting variable on each of the regressors in the model. This way we can measure the discriminatory capacity of the variables on a monthly or weekly basis using currently granted loans, closely following the market movement and providing early alerts if necessary.

2. Discriminatory Capacity of the Model: The model may also lose discriminatory capacity as the market changes. Unfortunately, this can only be determined exactly once a certain loan has been paid off or – in the opposite case – the respective customer defaults. We applied the following alternatives which can be considered in order to measure the model's performance periodically. We compute a score for each customer who defaulted or paid off the loan each month. Next a confusion matrix is constructed and percentages of false positives and false negatives are calculated. If any of these measures is above a predefined threshold, for example ten percent of the reported test accuracy, then an alert is given. This test allows following the performance of the model closely.

3. Change of Distribution of the Variables: The most challenging test, however, was to detect whether or not a significant change occurred in the distribution of the variables. The underlying question is if a small change in the variables' distribution is a risk for model performance. Any model should allow for certain variations in terms of the estimated coefficients, basically because a statistically correct sample of a population may also represent a slightly different one. Therefore standard K-S and χ^2-tests turned out not to be useful, since they are too sensitive to small variations. We solved this problem by constructing an empirical test as described in the following subsection using the model coefficients and their standard deviations as proxies. If a variable j has an associated estimated coefficient $\hat{\beta}_j$ and an estimated standard deviation $\hat{\sigma}_j$, it follows from the asymptotically normal behaviour of β coefficients in a logistic regression model (Hosmer & Lemeshow, 2000), that a 95% interval of confidence for the population parameter β_j is:

$$\hat{\beta}_j - 1.96\hat{\sigma}_j \leq \beta_j \leq \hat{\beta}_j + 1.96\hat{\sigma}_j \qquad j = 1, ..., N \tag{5}$$

The follow-up problem consists of measuring the shift between two different data sets: the original one used to construct the logistic regression model, and a second one with new cases. Formally, we have:

- Original data set $\mathbf{x}$ and estimated parameters $\hat{\beta}_j$ associated with variable j, $j = 1, ..., N$.
- New data set $\mathbf{x}'$, with new cases $\mathbf{x}'_i = (x'_{i1}, \ldots, x'_{iN})$, and observed outputs y'_i, $i = 1, \ldots, NC$ where NC is the number of new cases.
- Estimated default probabilities for the new cases $p(\mathbf{x}'_i)$ $(i = 1, ..., NC)$ obtained from the original models.

The respective literature provides some approximations to solve this particular problem. The closest model to the one presented here was presented by Zeira et al. (2005). It takes into account a measure of the shift, not just if a shift has occurred. The authors developed a statistical test for general models considering the output errors, assuming that they distribute normally, and that the variables are identically distributed. A variation of this approach, proposed by Cieslak & Chawla (2007), considers model evaluation in two steps: global discriminatory capacity of the model, and changes in the distribution of the variables instead of measuring them independently. Castermans et al. (2010) recently developed a framework for monitoring and validating the use of risk models within the context of Basel II accord. They consider a stability index to monitor the changes in the distribution of variables and a discrimination index for global performance.

4.2.1 Statistical Test for Model Follow-Up

The basic idea of the test we developed is to check whether the new coefficients $\hat{\beta}'_j$ estimated from the new dataset $(x'_{i1}, \ldots, x'_{iN}, y'_i) \in \Re^{(N+1)}$ $(i = 1, ..., NC)$ are still within the confidence interval belonging to the original estimators $\hat{\beta}_j$ of the variables.

Since the variables of the new dataset are constructed the same way as those from the original one (a necessary condition for applying the model), estimating the new coefficients $\hat{\beta}'_j$ is straightforward and can be done at a very low computational cost.

With the new estimators $\hat{\beta}'_j$, and their standard deviations $(\hat{\sigma}'_j)$ that are obtained from the new dataset, a new statistic for the population values of the variable β'_j can be constructed (Greene, 1993). If the new sample is large enough (necessary condition for estimating the logistic regression parameters), and considering the normality of the coefficients, the following is fulfilled:

$$\frac{\hat{\beta}'_j - \beta_{ref}}{\hat{\sigma}'_j} \rightsquigarrow t_{NC-N}, \tag{6}$$

where the statistic has a t-distribution with $NC - N$ degrees of freedom. The scalar β_{ref} represents the assumption about the population parameter. We constructed a statistical test to verify whether or not the estimated new parameters are still within the confidence intervals obtained for the original dataset considering two one-sided hypothesis tests, as shown in (7).

$$\begin{array}{c} H_0 : \hat{\beta}'_j = \beta_{lo} \\ H_1 : \hat{\beta}'_j < \beta_{lo} \end{array} \; and \; \begin{array}{c} H_0 : \hat{\beta}'_j = \beta_{up}, \\ H_1 : \hat{\beta}'_j > \beta_{up} \end{array} \tag{7}$$

where $\beta_{lo} := \hat{\beta}_j - 1.96\hat{\sigma}_j$ and $\beta_{up} := \hat{\beta}_j + 1.96\hat{\sigma}_j$.

The test checks whether or not the new parameters $\hat{\beta}'_j$ are still within the respective bounds determined by the previous model. In this case, i.e. the null hypothesis H_0 is not rejected, the distributions of the variables did not change significantly, which reconfirms the model's stability. To apply this test, the number of cases must be sufficient to ensure normal distribution of the beta coefficients, which should be given if the procedure is conducted every of three to six months.

Our approach differs from the introduced models in several points: first, by being model-dependent (uses the information and variability of the original model), it shows a more business-oriented focus on the concept drift approach. The other methodologies focus on whether a change occurred, or on the absolute strength of this change, not in if the change proposed is relevant for the model, as we believe should be. This fact ties the usefulness of our approach tightly to credit scoring, where logistic regression is widely used and the changes on variables can be much more sudden, in contrast with the other methods that present a more general approach. Another difference, contrasting Hellinger's distance for concept drift (Cieslak & Chawla, 2007) to our method, is that the former is only relevant for categorical variables, whereas our approach is applicable for both continuous and binary variables. The use of K-S and χ^2 measures, also common practice according to the respective literature, is not recommended, since they are known to be very sensitive to small changes in the variables' distributions, and are therefore prone to false alarms. Finally, a difference between our approach and the one by Castermans et al. (2010) is that the bounds on the "danger zone" of the latter are defined focusing on a percentage on the entropy in the comparison of the distributions, which is somewhat arbitrary, whereas our approach is tied to a hard bound on the limits given by the certainty on the original parameter estimation, which we believe is an advantage.

5 Results

In this section we present the results we obtained applying the proposed methodology to the governmental organization we worked with. The methodology employed, however, is generic and has also been used in our projects for the private sector. In particular we analysed two different datasets with loans that ranged from one to five years duration and amounts that varied between EUR 175 and EUR 17,500. Both datasets present an average granted loan of EUR 1,500.

The first data set (Universe U2) contains 41,200 long-term loans for new customers during a period of 12 years (1996 to 2007), and presents a high number of defaulters, with a default rate of 26.2%. The second dataset (Universe U4) contains 110,000 long-term loans for renewing customers during the same period, and shows a default rate of 17.9%.

The following subsections present results that were obtained in each one of the steps of the proposed methodology as described above, paying special attention to the results achieved by our newly developed steps for cut-off point construction and model follow-up.

5.1 Variable Selection

The repository created for datasets U2 and U4 initially contained more than 100 potential input variables. The goal of variable selection is to identify no more than 10-15 variables for model construction.

Variables commonly used in literature can be divided into three different groups: socio-demographic variables (customer provided, age, income, etc.), internal data (evolution of previous loans, other products, etc.), and external data (outstanding debts, checking accounts, etc.) (Anderson, 2007). All three were present in the original dataset, and all classical indicators of debt evolution were built, if possible. However, since the segment we analysed did not have access to financial services previously, the borrowers did not possess common debts and income variables. This fact made it even more difficult to develop the respective scoring systems than those presented in literature, and made most of "normal" variables useless.

Using simple filter methods for feature selection, we removed highly concentrated variables (i.e. more than 99% of cases that have the same feature value) and obviously irrelevant ones detected by K-S and χ^2-tests, reducing the number of remaining variables to fewer than 100. Subsequently, we applied forward selection and backward elimination for logistic regression obtaining a manageable number of features for each model.

During this variable selection process we maintained very close interaction with our customers, in particular with business experts and future users in order to assure suitable input variables for the respective models. This interaction is of utmost importance, since it adds business knowledge to the selection process and assures model acceptance by the intended users. However, in some cases we obtained surprising results, such as the case of the income variable, which was expected to be an important input. But as already mentioned, income seems not to be a very good variable when it comes to predicting micro-entrepreneurs' paying behaviour using credit scoring models. Our analyses confirmed this assumption. This is completely different from, for example, the mass consumer segment, in which the income varies much more, and is a relevant variable, especially when constructing indicators associating it with debt.

The final variables selected for U4 are divided into two groups. The socio-economic variables include Economic Activity (Activity), the sector of the economy that the customer is immersed in (through his/her job or company). The large number of activities was clustered to diminish the deviation and to improve

Table 2: Parameters Obtained for Logistic Regression model. New customers (left) and renewing customers (right).

(a) New Customers (U2)

Variable	β	S.E.	Sig.
Ownership_Owner	-.421	.044	.000
Ownership_Let	-.071	.057	.216
Ownership_Share	.184	.147	.210
LogAge	-.342	.047	.000
NumProp_One	.614	.058	.000
NumProp_More	.111	.066	.092
Activity_A	.320	.052	.000
Activity_B	-.022	.054	.677
Region_A	.093	.038	.015
Region_B	-.569	.044	.000

(b) Renewing Customers (U4)

Variable	β	S.E.	Sig.
Region_A	.092	.037	.011
Region_B	-.238	.032	.000
Ownership_Owner	-.293	.037	.000
Ownership_Let	.076	.051	.139
Ownership_Share	.354	.114	.002
NumProp_One	.497	.034	.000
NumProp_More	.196	.035	.000
LogAge	-.436	.046	.000
NumClosed	-.090	.004	.000
NumCurr	-.034	.013	.007
PrevArr	1.493	.032	.000
PercArr	.089	.018	.000
MaxArr	.001	.000	.000

interpretation of the variable, bringing the 47 different activities into three homogeneous groups of activities (Activity_A, Activity_B, and Activity_C); the ownership of housing (Ownership), whether the customer owns, rents, or has other types of agreements in his/her current home. Four classes are recognized: Owner, Tenant (Rent), Shares Tenant (Share), or others; the number of productive properties the borrower controls, i.e. properties that are necessary to develop the respective activities and therefore different from ownership, clusterized into zero, one, or more (NumProp_One, Numprop_More), ; a clusterization of the regions in the country (Region_A, Region_B, Region_C). Finally, the age of the customer in years, normalized to $[0, 1]$ (Age), or transformed using the natural logarithm (LogAge) is included. It is known that in general age should be treated as a discrete variable to account for non-linearities. However, in our case the range of ages was much more restricted than in normal consumer loans and the behavior was much more linear regarding age. For simplicity reasons we used this justification and treated age as linear variable.

The second group characterize the credit history of the customer, including the number of current or parallel loans (NumCurr), the number of closed loans (NumClosed), the average term for all loans granted previously (AvgTerm), a binary variable indicating whether the customer has been in arrears for any of his/her past loans(PrevArr), the percentage of the paid instalments of previous loans that have been in arrears (PercArr), and the maximum number of days that the customer was in arrears for any previous instalment (MaxArr).

5.2 Model Results

Applying the methodology for credit scoring described above, we obtained a logistic regression model for each one of the two datasets, U2 and U4. Table 2 displays the respective model parameters. The final models are characterized by 41,200 cases and 10 variables (in the case of U2), and 110,000 cases and 13 variables (in the case of U4). In both instances we used approximately 80% of cases for model construction and the remaining 20% for model evaluation.

To evaluate the obtained models, common accuracy measures were estimated for both datasets, with an unsurprising result: it is much more difficult to construct a logistic regression model for new customers than for renewing ones. The evaluation dataset associated with new customers (U2) presented an Area Under the Curve (AUC) of 0.6314, and a K-S maximum distance of 0.1991, while the evaluation dataset associated with renewing customers (U4) presented an AUC of 0.7795 with a K-S maximum distance of 0.4204. This represents a 15% increase in AUC and more than double K-S maximum difference between the two datasets. It is interesting to note that this result is very much in line with consumer credit scoring models, with similar adjustments to those observed in practice. The corresponding ROC curves are shown in Figure 1.

When studying the accuracy on a case-by-case basis, the models present 74.8% accuracy for non-defaulters, and 65.7% for defaulters from dataset U2, and 78.2% accuracy for non-defaulters, and 76.0% for defaulters from dataset U4, assuming a cut-off point of 0.5. Again, the overall measure accuracy provides better results for known customers (U4) than for new ones (U2). As can be observed from Table 3 and Table 4, however, the proposed models identify better among defaulters in U2, a result which is intuitive, since the behaviour of good payers is easily observed during repayment, and this additional information allows

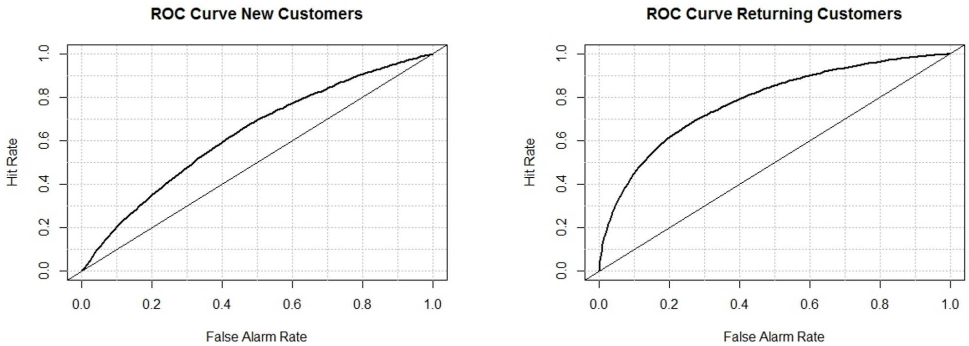

(a) New Customers

(b) Returning Customers

Figure 1: ROC Curves for the Two Datasets: New Customers (U2, AUC=0.6314) and Renewing Customers (U4, AUC=0.7795).

Table 3: Cut-off Table for New Customers (U2). Cut-off points maximizing accuracy (0.95) and minimizing total cost (0.55) are marked as **bold**, respectively. All monetary amounts in EUR thousands.

Cut-Off	Correctly Classified			Avg.	Avg.	Cost	Cost	Total
	Good	Defaulter	Total	Amount	Loss	(Good)	(Def.)	Cost
0.4	5,847	9,828	15,675	967	1,789	6,279	1,731	8,010
0.45	8,504	9,239	17,743	900	1,614	5,208	2,513	7,722
0.5	11,768	8,458	20,226	831	1,402	4,093	3,278	7,371
0.55	**15,340**	**7,457**	**22,797**	**784**	**1,226**	**3,120**	**4,092**	**7,213**
0.6	18,880	6,176	25,056	740	1,102	2,254	5,092	7,346
0.65	22,142	4,786	26,928	696	1,009	1,519	6,062	7,581
0.7	25,234	3,304	28,538	644	941	878	7,052	7,931
0.75	27,787	1,916	29,703	575	886	396	7,870	8,267
0.8	29,404	790	30,194	493	848	129	8,481	8,610
0.85	30,180	153	30,333	354	820	20	8,725	8,745
0.9	30,376	7	30,383	243	813	1	8,770	8,771
0.95	**30,396**	**0**	**30,396**	**0**	**812**	**0**	**8,770**	**8,770**
Total	**30,396**	**10,796**	**41,192**	-	-	-	-	-

for better discrimination of good payers. Furthermore, in our credit scoring projects we observed that good payers in general show a much more homogeneous behaviour than defaulters which present a multiplicity of reasons to default.

Analysing the obtained results reveals that the most important issue for determining paying behaviour is the way a customer handles his/her budget, which is reflected in variables associated with the number of days in arrears that the customer has.

5.3 Results of Cut-Off Point Construction

As was explained above, for cut-off point construction we need tables that consist of the solicited loans, the estimated default probability, and the loss incurred when defaulting. To determine cut-off points, the calculated default probability is replaced by the closest upper value of the 0.05 intervals used. That is, if the estimated probability is, for example, 0.543, it is replaced by 0.55. Segmenting the cut-off values allows grouping the loans and analysing the effects of changing the cut-off policy over a batch of loans, instead of on a per-loan basis, which greatly simplifies the analysis.

To construct the tables as shown e.g. in Table 3, we first identify the customers in each 0.05 interval and then calculate accuracy in the respective intervals. This is done separately for Good Customers (column 2 of each table) and for Defaulters (column 3 of each table). The total number of correctly classified cases is presented in column 4; see e.g. Table 3.

Table 4: Cut-off Table for Renewing Customers (U4). All monetary amounts in EUR thousands.

Cut-Off	**Correctly Classified**			**Avg.**	**Avg.**	**Cost**	**Cost**	**Total**
	Good	**Defaulter**	**Total**	**Amount**	**Loss**	**(Good)**	**(Def.)**	**Cost**
0.4	53,805	15,563	69,368	1,666	620	15,936	2,510	18,447
0.45	60,374	14,473	74,847	1,711	641	13,393	3,296	16,690
0.5	66,129	13,428	79,557	1,740	665	10,976	4,115	15,091
0.55	71,090	12,292	83,382	1,765	690	8,820	5,050	13,870
0.6	75,243	11,093	86,336	1,788	727	6,969	6,191	13,160
0.65	78,890	9,738	88,628	1,811	738	5,311	7,292	12,604
0.7	**82,004**	**8,312**	**90,316**	**1,876**	**746**	**3,957**	**8,432**	**12,390**
0.75	84,642	6,685	91,327	1,980	747	2,796	9,659	12,456
0.8	**86,737**	**5,054**	**91,791**	**2,118**	**766**	**1,818**	**11,158**	**12,977**
0.85	88,418	3,358	91,776	2,295	797	950	12,950	13,901
0.9	89,491	1,650	91,141	2,414	838	315	15,056	15,371
0.95	89,894	578	90,472	2,107	887	50	16,880	16,931
1	89,985	0	89,985	0	911	0	17,861	17,861
Total	**89,985**	**19,614**	**109,599**	-	-	-	-	-

The subsequent columns are used to estimate the related costs of rejecting a good applicant and accepting a bad applicant, respectively. Column 5 presents the average credit amount solicited by good payers who would be rejected if the corresponding cut-off point were used. Column 6 was calculated using Equation (2), considering the average observed loss per-customer up to that cut-off point.

Some considerations are relevant: first, collateral is not considered in this segment since the institution would incur in important social and direct costs when trying to recover such items. Consequently, there is no active policy for their collection. Second, the exposure of the loan considers only the amount due (amortization), not the interest collected. This is done to ensure transparency in the reported losses, since interest rates may vary from borrower to borrower.

Columns 7 and 8 present, for each possible cut-off point, the cost of rejected loans and the cost of accepted loans, respectively. The cost of rejected loans is obtained by multiplying the cost of rejecting a good applicant ($0.2644 \cdot$ Column 5) and the number of good applicants that would have been rejected for a given cut-off point (from Column 2, all good borrowers -last row- minus the value for the cut-off point). On the other hand, the cost of accepted loans is calculated by multiplying the cost of accepting a bad payer, from Column 6, by the number of bad payers that would have been accepted in that cut-off (from Column 3, all defaulters -last row- minus the value for the cut-off point).

The total cost (Column 9) is finally obtained by adding the cost of rejected loans and the cost of accepted loans (Column 7 + Column 8). Both resulting tables are presented in Table 3 and Table 4, respectively.

These tables reveal very interesting insights into the respective business. First, the proposed cut-off points are reasonable in the light of the risk associated with each of the two datasets. For new customers (U2), the total cost is the dominating factor taking into account that it is very difficult to determine the correct class of a first-time customer and that very high default rates have been obtained for this segment. Consequently, the suggestion is a very conservative 0.55 cut-off point, which translates into an acceptance rate (coverage) of only 45%. Evaluating the remaining requests carefully by a committee of experts considering the cut-off point associated with accuracy (0.95) is recommended.

The cut-off points for dataset U4 present a different picture, since for renewing customers much more information is available and the respective segment presents lower default rates. The cut-off point that minimizes total costs is 0.7 which translates into a direct acceptance rate of 85%. Loans with calculated reject probability greater than the cut-off point that maximizes accuracy (0.8) are rejected directly, leaving only a 7% of solicited loans to a committee.

It is also interesting to analyse the expected savings this policy would have generated. The current policy has a cost of 8,770,881 EUR for dataset U2, and 17,861,211 EUR for dataset U4 (considering only costs associated with loss). A conservative calculation can be performed to estimate the savings of using the models, considering that the cut-off point that minimizes costs is used, i.e. there is no committee.

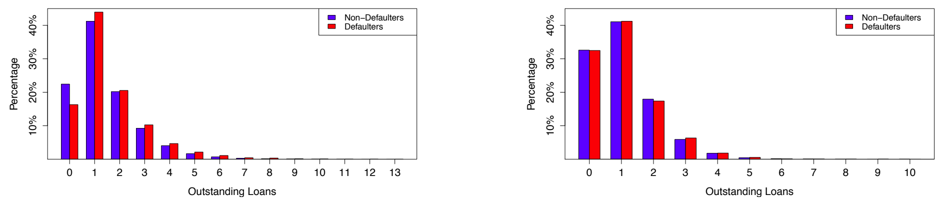

(a) 2000-2004

(b) 2005-2007

Figure 2: Distribution of variable NumCurr for U4, period 2000-2004 (left) and 2005-2007 (right)

Table 5: Follow-Up results, new customers (U2).

Variable	$\hat{\beta}'$	$\hat{\sigma}'$	$\hat{\beta}$	β_{lo}	β_{up}	t_{lo}	t_{up}	p_{lo}	p_{up}
Ownership_Owner	.283	.096	.559	.416	.701	-1.378	-4.347	.084	1.000
Ownership_Let	.216	.424	.527	.021	1.033	.461	-1.926	1.000	1.000
Ownership_Share	.167	.114	.367	.216	.518	-.426	-3.083	.335	1.000
LogAge	-.141	.119	-.659	-.821	-.496	5.707	2.984	1.000	**.001**
NumProp_One	-.295	.102	-1.108	-1.267	-.948	9.512	6.394	1.000	**.000**
NumProp_More	-.795	.224	-1.953	-2.310	-1.597	6.757	3.579	1.000	**.000**
Activity_A	.212	.090	.169	.035	.303	1.966	-1.003	1.000	1.000
Activity_B	.171	.101	-.546	-.693	-.400	8.519	5.632	1.000	**.000**
Region_A	-.167	.093	-.440	-.609	-.270	4.782	1.117	1.000	.132
Region_B	.015	.082	-.100	-.219	.019	2.860	-.041	1.000	1.000

Under this setting, the savings for our governmental institution, and therefore for society, when using the models in dataset U2 are $8,770,881 - 7,213,219 = 1,557,662$ EUR, which is equivalent to 4.2 EUR per loan granted. The savings in U4 are $17,861,211 - 12,390,124 = 5,471,087$, or a staggering 49.9 EUR per loan.

There is no need to emphasize the usefulness of the proposed methodology and of the use of credit scoring models in general, considering the substantial savings that are produced by the use of risk assessment models.

5.4 Follow-Up Results

In order to study the performance of the proposed follow-up methodology, we constructed models for datasets U2 and U4 using the loans granted during the period 2000-2004, and then we analysed the models' changes for the period 2005-2007. We divided the sample into an 80 percent training set for both periods, and a test set consisting of 20 percent of the data, constructing two test sets and two training sets. We trained four models, one for each period, and for each dataset (U2-U4). Figure 2 shows an example of how a particular variable (number of current loans from U4) changes its distribution in a period of three years, affecting its influence in the model.

To motivate the effects of concept drift we estimated the loss in discriminatory capacity applying the model with "old" data to the new test set: the AUC of the model built with new data corresponds to 0.6404, and to 0.5924 using the original model, representing a drop of eight percent in AUC. Finally, from the work of Blochlinger & Leippold (2006) it can be deduced that a drop of 0.01 in the K-S statistic can translate in a loss of up to two percent in the net utility of the lender, and in this case the statistic goes down from 0.2359 to 0.1622, representing 75 base points less, or up to 15 percent loss of utility for the lender, which can have catastrophic effects. We believe this reflects the strong need for proper model follow-up.

The results obtained for datasets U2 and U4 using the test proposed in Equation (7) are shown in Table 5 and Table 6, respectively. For each variable we computed the coefficients $\hat{\beta}'$ for the period 2005-2007, its standard deviation $\hat{\sigma}'$, the coefficients $\hat{\beta}$ for period 2000-2004, its lower bound β_{lo} and upper bound β_{up}, the t-statistics t_{lo} and t_{up} for both hypothesis tests presented in Equation 7 and the one-sided P values p_{lo} and p_{up}. These latter values can be interpreted as the probability that the statistic $\hat{\beta}'$ would differ as much as the boundaries obtained from the original dataset in the direction specified by the hypothesis just by chance, even though the parameter is actually within these boundaries (assuming that the null hypothesis is true). P values below 0.05 can be considered low enough to affirm that the new statistic differs significantly from the original one.

Table 6: Follow-Up results, renewing customers (U4).

Variable	$\hat{\beta}'$	$\hat{\sigma}'$	$\hat{\beta}$	β_{lo}	β_{up}	t_{lo}	t_{up}	p_{lo}	p_{up}
Region_A	.112	.078	.121	-.005	.247	1.492	-1.725	**.000**	1.000
Region_B	-.197	.075	-.135	-.273	.003	1.010	-2.649	1.000	1.000
Ownership_Owner	.268	.094	.689	.530	.848	-2.775	-6.154	**.003**	1.000
Ownership_Let	.379	.293	.919	.494	1.345	-.391	-3.294	.348	1.000
Ownership_Share	.122	.095	.632	.488	.775	-3.852	-6.874	**.000**	1.000
NumProp_One	.086	.074	-.687	-.808	-.566	12.045	8.790	1.000	**.000**
NumProp_More	-.018	.084	-1.143	-1.286	-1.000	15.047	11.654	1.000	**.000**
LogAge	-.397	.126	-.636	-.814	-.458	3.309	.488	1.000	.313
NumClosed	-.067	.007	-.114	-.131	-.096	9.536	4.307	1.000	**.000**
NumCurr	-.038	.035	.104	.057	.152	-2.659	-5.358	**.004**	1.000
PrevArr	1.351	.093	1.877	1.751	2.004	-4.322	-7.051	**.000**	1.000
PercArr	-.023	.026	.077	-.001	.154	-.837	-6.702	.201	1.000
MaxArr	.001	.000	.002	.002	.002	-3.832	-7.236	**.000**	1.000

It can be concluded from these experiments that the models present an important loss in performance, especially regarding prediction of non-defaulters. One of the main reasons for this change can be seen in the variables, where four out of ten present critical changes in their distribution for new customers (Age, Number of Properties in two levels, and Activity), affecting the performance of the model. The change is even greater for the universe of renewing customers, where nine out of thirteen variables present significant changes in comparison with the original data set. Furthermore, many of the affected variables become irrelevant in the new data set, which is a clear sign that the models need to be re-adjusted. These results underline the potential of the model follow-up methodology proposed in this paper.

6 Conclusions and Future Work

Granting loans to micro-entrepreneurs is a very important business opportunity in developing countries. As a country develops, granting such loans is slowly moving away from public institutions and is being considered a real business opportunity for private organizations, such as banks. The high risk, however, associated with micro-entrepreneurs is one of the main problems that hinders the expansion of this type of loans, and consequently a faster development of the respective economies. This explains the utmost importance of adequate risk assessment models that are tailored to the particular characteristics of micro-entrepreneurs.

In this paper, we have described the successful adaptations of the standard KDD process to the particular needs of public and private financial institutions that grant loans to micro-entrepreneurs, using logistic regression as the classification method. This procedure provides good results and a useful interpretation of the respective input variables.

Traditional credit scoring variables such as "income" were not relevant. Micro-entrepreneurs have similar incomes and therefore this variable does not discriminate. We also tried related variables using income, such as income/debt, etc. which neither discriminated well.

Another important conclusion is to use variables that can be corroborated. Ambiguous variables describing characteristics that cannot be proven easily are not useful since customers and/or salespersons know how to answer certain questions in order to improve their chances of getting the loan.

The proposed cut-off point methodology explicitly quantifies the lost benefit for society when a loan for a good customer is not granted. Our numerical experiments underline the potential impact such a methodology might generate, and help in quantifying the benefit of using statistical models in practice, with important savings when compared to the absence of such models.

Micro-entrepreneurs represent a very volatile market, and are therefore very sensitive to changes in economic conditions, which makes the nature of the respective credit granting operations very dynamic and subject to constant change. This is reflected by shifting risk factors. Tools that detect such shifts are very attractive for practitioners but not always available in standard solutions for credit scoring. In this paper we introduced statistical tests for model follow-up that were developed in our projects and provided excellent results.

By using adequate methodologies for credit risk management, the market of loans for micro-entrepreneurs will continue to grow as more private companies will offer such loans; see e.g. Kim & Sohn (2007) and Kim & Sohn (2010) for a similar situation. This, in turn, will foster their capacity to innovate and generate growth. OR-methodologies will contribute to a sustainable development of the respective countries as has been shown already for many other cases in White et al. (2011).

Acknowledgement

The first author acknowledges CONICYT for the grants that supports this work (AT-24110006, NAC-DOC: 21090573). All authors acknowledge the support of the institution which provided the data. The work reported in this paper has been partially funded by the Institute of Complex Engineering Systems (ICM: P-05-004-F, CONICYT: FBO16) and the Finance Center of the Department of Industrial Engineering, Universidad de Chile, with the support of bank Bci.

References

Anderson, R. (2007). *The Credit Scoring Toolkit*. Oxford University Press.

Arrow, K. J., & Debreu, G. (2002). *Landmark Papers in General Equilibrium Theory, Social Choice and Welfare*. Edward Elgar Publishing.

Baesens, B., Van Gestel, T., Viaene, S., Stepanova, M., Suykens, J., & Vanthienen, J. (2003). Benchmarking state-of-the-art classification algorithms for credit scoring. *The Journal of the Operational Research Society*, *54*, 627–636.

Bellotti, T., & Crook, J. (2008). Credit scoring with macroeconomic variables using survival analysis. *Journal of the Operational Research Society*, *60*, 1699–1707.

Blochlinger, A., & Leippold, M. (2006). Economic benefit of powerful credit scoring. *Journal of Banking & Finance*, *30*, 851–873.

Brown, I., & Mues, C. (2012). An experimental comparison of classification algorithms for imbalanced credit scoring data sets. *Expert Systems with Applications*, *39*, 3446–3453.

Castermans, G., Hamers, B., Van Gestel, T., & Baesens, B. (2010). An overview and framework for PD backtesting and benchmarking. *The Journal of the Operational Research Society*, *61*, 359–373.

Cieslak, D., & Chawla, N. (2007). Detecting fractures in classifier performance. In *Proceedings of the Seventh IEEE International Conference on Data Mining* (pp. 123–132). Department of Computer Science and Engineering, University of Notredame.

Instituto Nacional de Estadísticas, I. (2002). Primera encuesta de las micro, pequeñas y medianas empresas [first survey of micro, small, and medium companies]. Retrieved online 14 February, 2012.

Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). The KDD process for extracting useful knowledge from volumes of data. *Communications of the ACM*, *39*, 27–34.

Finlay, S. (2011). Multiple classifier architectures and their application to credit risk assessment. *European Journal of Operational Research*, *210*, 368–378.

Greene, W. H. (1993). *Econometric Analysis*. Prentice Hall.

Hand, D., & Henley, W. (1997). Statistical classification methods in consumer credit scoring: a review. *Journal of the Royal Statistical Society Association*, *160*, 523–541.

Hosmer, D., & Lemeshow, H. (2000). *Applied Logistic Regression*. John Wiley & Sons.

Kim, H. S., & Sohn, S. Y. (2007). Random effects logistic regression model for default prediction of technology credit garantee fund. *European Journal of Operational Research*, *183*, 472–478.

Kim, H. S., & Sohn, S. Y. (2010). Support vector machines for default prediction of SMEs based on technology credit. *European Journal of Operational Research*, *201*, 838–846.

División Empresas de Menor Tamaño, M. d. E. (2009). Encuesta longitudinal de empresas [longitudinal survey of companies]. Retrieved online 14 February, 2012.

Ministerio de Economía (2012). Segunda encuesta de microemprendimiento [second micro-entrepreneurship survey]. Retrieved online 5 May, 2012.

Ozdemir, B., & Miu, P. (2009). *Basel II Implementation*. McGraw-Hill.

Schreiner, M. (2000). Credit scoring for microfinance - can it work? *Journal of Microfinance*, *2*, 105–119.

Setiono, R., Baesens, B., & Mues, C. (2009). A note on knowledge discovery using neural networks and its application to credit card screening. *European Journal of Operational Research*, *192*, 326–332.

Superintendencia de Bancos e Instituciones Financieras (2008). *Compendio de Normas Contables [Compendium of Accounting Rules]*. SBIF.

Thomas, L., Crook, J., & Edelman, D. (2002). *Credit Scoring and its Applications*. SIAM.

Tong, E. N. C., Mues, C., & Thomas, L. C. (2012). Mixture cure models: if and when borrowers default. *European Journal of Operations Research*, *218*, 132–139.

Van Gool, J., Verbeke, W., Sercu, P., & Baesens, B. (2011). Credit scoring for microfinance: is it worth it? *International Journal of Finance & Economics*, *Available Online*, 24 January, 2011.

White, L., Smith, H., & Currie, C. (2011). OR in developing countries: A review. *European Journal of Operational Research*, *208*, 1–11.

Zeira, G., Last, M., & Maimon, O. (2005). Advanced techniques in knowledge discovery and data mining. chapter Segmentation on Continuous Data Streams Based on a Change Detection Methodology. (pp. 103–126). Springer.