Sequential Banking Externalities*

Giacomo De Giorgi

GSEM-University of Geneva

BREAD, CEPR, and IPA

Andres Drenik

Columbia University

Enrique Seira

ITAM

Stanford University's Hoover Institution

Abstract

The ability to borrow sequentially from multiple lenders is a standard feature of credit markets that theoretically may lead to high default and inefficiency, yet little is known about its prevalence in practice and the risks it induces for the financial sector. Using data on all loans in Mexico we first show that sequential banking is pervasive. Second, using a regression discontinuity design we show it causes a 92% increase in default on *sequentially prior* cards and 48% for non-card loans, resulting in average losses of 18% of total debt, an important externality on previous lenders. Third, we find that additional credit induces default only for borrowers in the "bancarization" margin, those with lower scores, and not on higher scores customers. These results confirm that financial inclusion is hard, and might explain why the previous literature has been inconclusive on this issue. Our results have important implications for the design of universal-default clauses recently implemented in the US and Mexico.

Keywords: Sequential banking, Externalities, Default, Financial inclusion, Universal default clauses
JEL: D14, E51, G21

*We thank Bernardo Garcia Bulle and Eduardo Laguna for excellent research assistance. We also want to thank Liran Einav, Andres Liberman, Brigitte Madrian, Neale Mahoney, Jean-Charles Rochet, Johannes Stroebel, Jonathan Zinman for helpful comments. Enrique Seria thanks the Mexican Central Bank where part of this resesearch took place during his stay in the reserch department. Giacomo De Giorgi acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075) and ECO2011-28822, the EU through the Marie Curie CIG grant FP7-631510. Authors' contact information: giacomo.degiorgi@gmail.com.

1 Introduction

A standard feature of consumer credit markets is that borrowers can apply sequentially for loans from different lenders. However, approval decisions of *posterior* lenders may affect the profitability and risk of previously opened loans, and therefore impose an externality on *prior* lenders. Although the sequential banking phenomenon is anecdotally prevalent, the implications and relevance of this arrangement have not been studied empirically.[1] Bizer and DeMarzo (1992) study the issue theoretically in a moral hazard model where more debt leads to lower repayment effort, and conclude that compared to a one-lender model, a "sequential banking equilibrium" involves higher interest rates, more borrowing, higher default, and inefficiency, even when lenders anticipate that the borrower will ask for loans sequentially.

The first contribution of this paper is to establish the actual prevalence of sequential banking in Mexico using a random sample of 1 million borrowers, representative of *all* formal borrowers in the country (approximately 57 millions). We find that the majority of borrowers have more than 1 loan simultaneously, with one third having three or more loans. The time in between getting loans is short: 12 months on average between the third and fourth loan, but only 7 months between the sixth and the seventh loan.

Second, we study whether sequential banking generates a risk for the financial sector using loan *application* data from one of Mexico's largest banks –Bank A henceforth– that serves costumers across the income distribution. This bank has more than 1,000 branches covering all 32 States. This by itself is interesting as a large part of the previous literature focused on smaller lenders typically serving low income customers. More importantly, we are able to overcome two crucial limitations that have limited the study of sequential banking.

[1] In Mexico and in the US, consumers with at least one credit card have more than three cards on average for which they have presumably applied sequentially, on top of other loans. Cappelletti et al. (2017) highlight how little we know about the effects of multiple banking on default risk. The determinants of the number of bank relationships and the advantages and disadvantages of multiple relationships have been studied in the literature (e.g. Ongena and Smith (2000), Detragiache et al. (2000), Brunner and Krahnen (2008), Gobbi and Sette (2013), Bris and Welch (2005)). A recent working paper by Axelson and Makarov (2017) studies in a theoretical setting the consequences of sequential credit market searches for investment financing. Those markets can lead to substantial inefficiencies, which a credit bureau is not able to eliminate.

One of them is the scarcity of rigorous empirical designs that enables researchers to make causal claims about the effect of loan approval on default in *pre-existing* loans. The other is the necessity of observing default behavior on *all* the loans held by an individual and the order in which they are obtained.

We have the luxury of observing all loan applications (both rejected and accepted) to Bank A, and of exploiting its approval rules in a regression discontinuity design (RDD). This allows us to establish a causal relationship between loan approval and outcomes. At the same time, from the Credit Bureau we obtained data on the universe of formal loans for the individuals who were applicants to Bank A. Therefore, we are able to study default and delinquencies on previous and new loans. We are also able to study default by type of loans, and estimate whether getting a credit card increases or decreases default in personal loans. Our results indicate that externalities in default due to sequential banking can be substantial, leading to a *doubling* of default rates on new and prior loans. This result verifies the main assumption of theoretical papers such as Bizer and DeMarzo (1992), and means that sequential banking could lead to inefficiently high interest rates and risk in financial markets in equilibrium. This has implications for current regulation. For instance, current no-universal-default clauses that forbid banks to adjust prices and terms as a function of default in other banks could hinder bank's ability to price away the externality.

The third contribution of the paper is to study an actual financial inclusion effort at the country level and in a low inclusion country. Many international institutions such as the World Bank encourage vigorous efforts to give financial services to underserved (extra-marginal) consumers, but there is little evidence on the effects of such efforts. Our analysis benefit in this respect by that fact that during our sample period Bank A tried to incorporate extra-marginal borrowers (those with lower scores) in a notable financial inclusion effort. In doing this, it experimented with different approval thresholds, this experimentation allows us to estimate the elasticity of default to credit along the credit score distribution. We can therefore document some of the difficulties of financial inclusion: credit-default elasticities were much higher for extra-marginal borrowers, to the point that Bank A backed away from its financial inclusion effort as it found it unprofitable.

A fourth and related contribution is to rigorously document that the credit-default elasticity varies substantially by credit score. While low score applicants respond with substantially more default to the award of a new loan as in moral hazard models, high score borrowers actually *decrease* default (as in Dobbie and Skiba (2013)), suggesting that their use of "extra" credit is potentially tied only to the necessity of smoothing temporary shocks. This heterogeneity may be a key to explain why different authors (e.g. Karlan and Zinman (2009), Adams et al. (2009), and Dobbie and Skiba (2013)) have reached different conclusions about the existence and significance of moral hazard in consumer loan markets; pooling different populations could mute the elasticity. Although we cannot pin down the exact mechanism for the difference between high vs low score consumers, we find that low score borrowers have a larger propensity to borrow, accumulating twice the amount of debt than high score borrowers as a response to the same increase in the credit line, suggesting that they have a higher need for liquidity (as in a typical case of a credit constrained population), while high score applicants may apply for the card to have a line available for precautionary motives, i.e. a source of funds for the rainy days.

The paper proceeds as follows. We first assess the validity of the RDD strategy by showing that applicants from the left and right of the thresholds (in terms of credit scores) are statistically "identical" in observable characteristics. We then show that the probability density function of the applicants' credit score is smooth at the two thresholds, indicating that applicants are unable to manipulate their credit score at the thresholds. Second, we show that Bank A's treatment is large and persistent: the probability of credit card approval increases by about 45 percentage points (pp) just to the right of the score threshold, over a basis of 3pp to the left. This jump in the probability of approval translates into an immediate

47 percent increase in the total credit limit (on average) or about 16,000 MXN (960 USD). The difference in the number of cards between the control and treatment groups persists for more than two years.

Finally, we present our main results using this exogenous variation in the amount of credit. We find that for applicants close to the 670 cutoff, getting Bank A's card increases the cumulative probability of default on all credit cards (previously held and new) by 25pp in the next 18 months compared to a control group mean of 23pp.[2] In other words, we find that an extra 1,000 MXN (60 USD) is associated with a 1.5pp additional probability of default for the lowest score group. In comparison, we find a small (and less precise) *negative* effect on default for those applicants close to the 700 threshold. Focusing on externalities, i.e. on cards *already existing at application*, we estimate an average effect on default of 18pp over a base of 20pp for the applicants close to the 670 cutoff. We also find a spillover effect on other types of loans (like auto loans, personal loans, etc.), what we call extensive margin externalities. For this set of loans, we find an increase in the cumulative probability of default of 20pp over a control base of 35pp. These magnitudes forewarn that these externalities are large and merit careful consideration.

Following the recent literature, one can link these results with economic theories of asymmetric information. Note that our empirical strategy implicitly does three things. It holds constant the selection of applicants at the threshold, gives more loans to some of them quasi-randomly, and then measures the effect of this on default, documenting a strong positive effect for lower score individuals. Several papers take this result as evidence of the existence of asymmetric information, since a large class models of asymmetric information generate this positive correlation between default and loan quantities (or prices). In particular, when holding selection constant as in Karlan and Zinman (2009) and Adams et al. (2009), these papers take the positive relation as evidence of moral hazard. Our methodology supports

[2] This is the treatment on the treated, the intent to treat effect is 12pp. By cumulative probability of default, we mean that the credit card was at least 3 months delinquent during any time between application and 18 months after.

the same interpretation, but since we focus on the debt-default externality itself and less on the model primitives behind it, we do not stress the moral hazard interpretation in the paper.

While we are the first to estimate default externalities in sequential banking, several papers have estimated default elasticities in different contexts. Karlan and Zinman (2009) focus on microcredit costumers in South Africa and, by randomly varying current and future interest rates, they document that a 100 basis points decrease in the promised future interest rate causes a decrease of 4% in default, but they find little effect to changes in current interest rates, suggesting low moral hazard in this dimension. The paper by Adams et al. (2009) is similar to ours in that they focus on variation in the quantity of credit instead of its price. They study the subprime auto loans market in the US and show that, conditional on selection, an increase of $1,000 USD in the size of auto car loans leads to a 16% higher hazard rate of default. They also look at variations in current prices and, in contrast to Karlan and Zinman (2009), find large moral hazard effects. Interestingly, using an RD design in the context of payday lending in the US, Dobbie and Skiba (2013) find that a $50 USD larger loan leads to a 17 to 33 percent *decrease* in default. One interpretation they provide is that greater liquidity may allow for better debt management and timely payment.

Our paper differs from the ones mentioned above in several important dimensions. First, while all the above papers study the *subprime* market, we study a market for middle income individuals (applicants are mostly located in the second to fourth quartile of the income distribution). Second, the size of the treatment (i.e. approved loan quantity) is bigger than most treatments studied in this literature, which gives us substantial statistical power and perhaps explains the large effects we find. Third, we are able to look not only at the extensive margin of default, such as the probability of default, but also at the intensive margin, i.e. the number of defaulted lines across types of loans. Fourth, we can study heterogeneity of the response at different levels of the risk score and show widely different responses, with much larger default effects for lower score applicants. Fifth and most important, we estimate default *externalities* with a causal design.

In a recent concurrent paper, Agarwal et al. (2017) study the marginal propensity to lend by banks and to borrow by credit card holders in the US when credit limits increase due to a fall in the banks' cost of funds. They document substantial heterogeneity in the marginal propensity to borrow by consumers; and when discussing the marginal propensity to lend, they document that it is less profitable for banks to increase limits to existing clients with lower FICO scores. We focus instead on how new loans (extensive margin) lead to defaulting in sequentially prior loans. While the unit of analysis of Agarwal et al. (2017) is the credit card and they look at borrowing there, ours is the individual and we use outcomes for all types of loans. Agarwal et al. (2017) results are very important for stimulus policies designed to increase credit use, ours are helpful at quantifying the risks and externalities generated by sequential banking and in the design and costs of no-universal-default clauses.[3][4]

The rest of the paper proceeds as follows: Section 2 describes the institutional features of the market we study and the data used in the analysis; Section 3 presents the empirical strategy while Section 4 shows our main results. Section 5 presents robustness checks. Finally, Section 6 concludes.

2 Context and Data

2.1 Some Background

The Mexican credit card market is relatively underdeveloped and concentrated. The five largest banks held a steady market share of close to 90% from 2001 to 2012 in terms of the number of cards and outstanding credit card debt. Compared to the US, which has

[3]No-Universal-Default clauses make it illegal for banks to cancel a loan or increase its interest rate as a function of the client's behavior in servicing other loans.

[4]Another related paper, Lieberman (2016), studies debt renegotiation and shows that erasing bad credit histories generates more lending and higher default with other banks. Thus, he shows that credit history information is very valuable. We keep credit history reporting constant and show that, in spite of the incentives to limit default that credit history creates, sequential banking externalities exist and are large.

thousands of credit card issuers, Mexico has only about 20 (in Mexico only banks are card issuers). Average credit card interest rates have been close to 29 percent per year, while the government federal discount rate (TIIE) has remained between 5 and 7 percent (Banxico (2013)).

Mexico also has a relatively low penetration of cards, owing perhaps to a history of nationalization, privatization and recurrent financial crises in the 1980s and 1990s, including the Tequila crisis of 1994. Even in 2004, ten years after this crisis, there were 0.13 credit cards per person in the country compared to 0.35 in Argentina and 0.38 in Brazil (US (2008)). As of the early 2010s, the coverage rate was still low: there were 30 cards per every 100 inhabitants, whereas the analogous number for the US was 120 in that same year.[5] Low penetration is not only a feature of the credit card market in Mexico, in fact total credit to the private sector over GDP is close to 30% only, whereas for developed countries it is often above 100%.

A fraction of the penetration gap was closed during a high growth period between 2002 and 2008, in which the number of cards awarded grew at a rate of 9.9 percent per year. For the purpose of this paper it is important to note that this growth came in no small way from banks issuing new cards to *existing* cardholders. In 2008, 41% of new cards went to people who already had cards.[6] In fact, between 2006 and 2008 the number of cards held by the average cardholder increased from 3.4 to 4.2 (Banxico (2009)). This is reflected in the distribution of the stock of cards in the economy: in 2010 half the cardholders had one credit card, while 20%, 11% , 7%, 12% had two, three, four and five or more credit cards. Awarding cards or loans to borrowers that already have cards or loans is even more common in the US, in particular above 90% of new cards go to people who already have at least 1 card. In the case of Mexico, this increase in the number of cards was accompanied –although we do not claim causality here– by increases in default rates: while the non-performing card

[5]See Comision Nacional Bancaria y de Valores (2013) and Federal Reserve Bank of New York (2010).
[6]This number was 45% in 2007 and 41% in 2008.

debt was 4.9% as a percentage of total credit card debt in 2002, it was 12.2% in 2012. Part of the increase may be due to the incorporation of riskier marginal borrowers, while another part to awarding cards to borrowers that already had cards and substantial debt.

Sequential banking is a real and prevalent phenomenon. Using a random sample of 1 million borrowers in Mexico, Figure 1 shows that, conditional on having an active loan, only 47% of people have a single loan, while the rest have several loans. A non-negligible fraction have 5 or more loans outstanding. Additionally, people take 28 months to get a second loan after getting the first on average, but time shortens significantly thereafter, with less than 9 months between their fourth and fifth loan for instance.

Although in this paper we do not study the specific reason why borrowers default, it is worth going over the costs and benefits of this decision. The main benefit of default for a borrower is of course not paying the debt owed. After a default episode, Bank A and most banks in Mexico do not go after debts smaller than 60,000-100,000 MXN, as collection costs are high and courts slow and ineffective. When faced with credit card default, banks in Mexico sell the defaulted debt to collection agencies at about 90% discount. Thus, defaults are highly costly for banks. On the other hand, the main cost of default a borrower faces is a negative credit history at the credit bureau. Castellanos et al. (2017) have found that a loan default in Mexico subtracts close to 100 points from credit scores and makes it much harder to get loans in the future.

Interestingly, in Mexico it is illegal for banks to cancel a loan or increase its interest rate as a function of the client's behavior in servicing *other* loans. The authority considers "universal default" clauses abusive.[7] The regulation states that "Abusive clauses include those that... (g) permit the modification...of what was agreed in the contract without the consent of the user, unless it is in the benefit of the latter." Central bank regulators told us in correspondence that they do not know of any credit contract in Mexico that allows default in one contract to affect the conditions of another, in compliance with the regulation.

[7] See http://e-portalif.condusef.gob.mx/reca/manual/DCG_cla_abu.pdf

We also looked at evidence using our data; we estimated a regression of loan closings as a function of default at other banks and we found zero correlation. The US has also undergone a regulatory push back against universal default contract clauses culminating in the Credit Card Act of 2009 where most forms of the practice were outlawed. The Credit Card Act of 2009 limited "universal default" and prohibited retroactively increasing interest rates on existing balances as a function of behavior with other lenders. We take no position on this issue here, but we want to highlight its importance for our purpose, as it limits what banks can do to mitigate sequential banking externalities.[8]

2.2 Data, Approval Decision, and Financial Inclusion Effort

Our empirical analysis relies on two main data sources. The first one comes from all credit card applications made to Bank A between January 2010 and April 2012, from applicants that were not clients of Bank A. Bank A has more than 1,000 branches covering all 32 States of Mexico, making this not a study of a niche lender in a particular location, but a country wide phenomenon.[9] An observation consists of an individual credit card application. The data contain all the information recorded by Bank A at the moment of the application and used during the approval decision, including information on applicants' credit score, date of application, self-reported annual income, gender, and type of credit card applied for. It also contains some credit bureau information at the application date from which we can deduce the number and size of credit lines at application, as well as an identification number that allowed us to merge application data with credit bureau data. The data also include the bank's approval decision, type of card, interest rate and credit limit awarded in case of approval.

We merged Bank A's data with a second dataset coming from the Credit Bureau (CB)

[8] http://www.ausubel.com/creditcard-papers/ausubel-testimony-12february2009.pdf argues that penalties for default in other banks were much higher than the increased risk this represented.

[9] The reason for excluding applicants who had a savings or checking account at Bank A is that Bank A uses a very different approval process for them. Their process relies little on the credit score, and therefore there is no discontinuity in the approval decision that can be exploited in our empirical design.

that contains *the universe of loans* for all the applicants, closed and active ones. We have two snapshots of these data, one from January 2010 and the other from June 2013. The first snapshot occurs before our sample period begins and we use it to run tests of balance of pre-treatment characteristics. We use the June 2013 snapshot to measure outcomes. In this second dataset an observation corresponds to a single credit line. For each line we observe its type (mortgage, personal loan, credit card, etc.), opening and closing date, the credit limit and debt at the time the snapshot was taken, the current status of the credit (late payments, default, etc.) and the monthly payment history up to the last 6 years (however, we do not observe the interest rates of each credit). Such a dataset allows us to precisely measure all possible delinquencies for every month. Importantly we can look into effects 18 months after treatment for all applicants (and up to 3.5 years for the early ones). A card is delinquent one month if the minimum payment corresponding to that month was not paid. In keeping with the legal definition in Mexico, and with the literature (e.g. Gross and Souleles (2002)), a card is considered to be in default if it is delinquent for 3 consecutive months or more.

The CB data has strengths and weaknesses. Variables like dates of opening and closing of loans, and on the history of delinquency, are actively verified by authorities, banks and consumers themselves. Default is closely monitored by authorities and banks, since reserve requirements depend on it. Consumers also pay attention since bad credit history affects their access to credit. Consumers have the right to ask for amendments if information about them is not correct. The variables regarding limit and debt in the CB are subject to less verification. Moreover, while we have monthly information for default, for debt and limit we only observe their status at the time of the snapshots we obtained. This gives rise to two problems. Since card applications were filed in different months, a different number of months elapsed for different applicants in the 2013 snapshot. Also, debt is highly variable, potentially changing on a daily basis. Since we only have information from a single snapshot, this noise reduces the power of statistical tests. For these reasons, we focus mostly on default instead of debt or credit limit. We use a third dataset, from the social security administration, in order to have a verifiable measure of formal sector income.

A crucial variable in our analysis is the applicant's credit score at the moment of application, since it is our running variable in the research design. The credit score is computed by the Credit Bureau and sold to Bank A. We use exactly the same score that Bank A used for its approval decision. The calculation of the score is similar to the ones in the US (in fact the scoring method was designed by Fair Isaac, the leading credit scoring company in the US), using the individual's credit history, types and number of credits in use, among others. Although the exact formula is proprietary, credit scores are calculated using prior credit behavior and do not use any information about the individual's occupation, income, employment history, gender, age, or geographic location. Because we rely on an RD design, we use data from applicants that are within a 30 score-point range around the respective thresholds, i.e. with scores between 640 and 730. Of the total pool of applicants to Bank A, 46% fall in this range.[10] 34% of applicants have scores below 640, 44% below 670 (lowest cutoff), 60% below 700 (highest cutoff) and 80% below 730.

For our empirical design to be valid, we need that consumers are unable to manipulate their credit score with precision around the threshold. This is indeed the case for several reasons. First, Bank A's credit score threshold policy is not communicated to loan officers at the branches. Loan officers input the loan application information into the computing system and the system gives back an approval or rejection decision at this first appraisal stage. This happens without loan officers knowing the reason for rejection nor the cutoff that is currently used at that moment. Second, nobody knows the formula generating credit scores at the credit bureau, not even Bank A. This makes it impossible to manipulate the formula with precision. Third, the formula uses the whole credit history and therefore operates with a significant lag. Thus, it may take months or years to change the score significantly. Finally,

[10] Although not straight-forwardly comparable, in the US about 20-25% of those with a Vantage 3 credit score (provided by Experian) would have scores between those two extremes.

anecdotally few people know their credit scores in Mexico. Keys et al. (2010) find that a similar RD strategy is valid in the US. Although this narrative is persuasive, in Section 3.1 we present statistical evidence showing that there is no manipulation of the running variable.

Bank A's card approval policy proceeds as follows. If the applicant does not have a credit history in the credit bureau he is immediately rejected, otherwise he has to pass the current credit score threshold in the computing system at the moment of application. Applicants with scores lower than this threshold are almost always rejected (only about 3% of cases override this rule), while applicants with score above the threshold pass to a second credit appraisal stage. In the second stage, other variables (like income and time of oldest account) come into play, and the application may be rejected in this second stage based on these other variables. Notice that during this approval process the interest rate is not tailored to each applicant. Instead, interest rates are specific to the type of credit card obtained by the applicant (classic, gold, platinum and infinite). In our sample, the vast majority of applicants (92%) applied for the gold credit card. Out of the approved applications, 79% of applicants received the gold card and 17% got the classic card (which means that some gold applications were downgraded). The type of card that applicants end up receiving depends on other discontinuities further down the application process (thus, neither the applicant nor the bank employee has any influence on the final outcome). Therefore, interest rates are specific to the type of credit card and homogeneous across cards at any given point in time. For example, in April 2010, the gold credit card had an interest of 32% per annum, while the classic card had an interest rate of 48% per annum.

Importantly, Bank A changed the threshold twice during our sample period. The cutoff it had always used was 700 points, and most observations in our data come from periods where this threshold was in place, January 2010 to April 2011. However, the bank made an effort to serve extra-marginal consumers that had scores lower than these. Therefore, between June and November of 2011, it lowered its threshold to 670, see Figure 2. Finally, Bank

A increased the threshold to 680 between December 2011 and April 2012.[11] These changes were blind to the loan officers at bank branches and the length of time it would be active was unknown since it would depend on realized profitability. Other than changing the threshold, everything else in the approval process remained constant: the bank offered the same interest rates, same card products and conditions. Fortunately for us, this constitutes a window into a financial inclusion effort rarely studied in the literature. Since the vast majority of the point estimates obtained for the 680 sample lie in between those obtained for the 670 and 700 samples, we decided not to report them here. Ultimately, these 3 thresholds allow us to identify the parameters of interest in three meaningful parts of the credit score distribution. Figure 3 shows that this 30-point difference represents a significant and meaningful financial inclusion effort: while people with a credit score of 700 have a probability of default of around 4pp in the next 12 months, this probability increases to about 6pp for those with a credit score of 670 (a similar relation applies the US population).

2.3 Descriptive Statistics

Panels A, B and C of Table 1 show pre-treatment summary statistics using data from Bank A collected at the moment of application and from the Credit Bureau's December 2010 snapshot. We provide statistics for the pooled sample of applicants, as well as by credit score threshold using a symmetric interval of 10 points centered around the respective threshold. In the description of the table, we refer to applicants in the [665,675] interval as the 670 score applicants, and to those in the [695,705] interval as the 700 score applicants.

We want to highlight a subset of statistics, starting with monthly income as reported to the social security administration. Income varies with the score: it is 11,055 MXN (about 660 USD) for the 670 applicants and 14,199 MXN for the 700 applicants. This means that when we talk about going after extra-marginal borrowers by offering loans to lower credit

[11] We discard all applications made in May 2011 because Bank A was experimenting with two simultaneous cutoffs, which made the discontinuities in the probability of approval very small. We also discard a very small number of observations where the same person applied more than once to Bank A.

score applicants, it also means giving loans to lower income applicants.[12] This level of income would place our applicants' sample in the third quarter of the household income distribution in Mexico (INEGI (2012)). However, given the large variation in income, applicants kept in our estimation sample span a large portion of the Mexican income distribution, with most of the observations concentrated between the 5th to 8th higher deciles. From the CB data, we see that the population in the study has on average been in the credit bureau records for almost 8 years and has an average of 3.7 loans – these include personal loans, car loans, mortgages, credit cards, etc. Applicants in the 700 group have 37,038 MXN pesos in total outstanding debt, while those in the 670 set have 30,043 MXN. This means that our applicants use loans other than cards since the average credit card debt is 8,439 MXN (about 505 USD, not reported), about a quarter of total debt.

Our measures of delinquency and default are defined at the applicant (not the credit) level. For Table 1 (pre-treatment) we define the probability of delinquency in credit cards as equal to one if the person has had *any* credit card with 60 to 90 days past due *at any point in time* from the earliest month with available information of the card to the date of application to Bank A.[13] Note that we are using a cumulative measure of delinquency and not measuring delinquency at a specific point in time. We do this because default may lead to the closing of the loan, and we want to consider a loan as defaulted even if it is closed by the 2013 snapshot.[14] The probability of default is analogously defined, but considering loans that were 90 days or more past due. This corresponds to the standard definition of default used by the Mexican authorities (and has legal consequences in Mexico in terms of

[12] We were able to merge the applicants sample with administrative data from the social security. Although given the high degree of informal jobs and the quality of the matching variable, we could only match 21% of them. We also observe self reported income of all applicants filed with the application, but we do not use it here as the bank does not verify it and we think is over-reported and noisy; it tended to be higher than the income reported by employers to the social security for the applicants we could match. On average it was 27,350 MXN (about 1,640 USD) per month (unreported in Table).

[13] In Section 4.2 we will measure cumulative default from the time of application instead.

[14] A separate issue is that the CB by law has to delete defaulted loans from their dataset after some years as a function of default severity. If the defaulted debt is less than 113 MXN the bad credit history is deleted within a year, if it is between 113 and 2260 MXN it is deleted after 2 years, those between 2260 and 4520 MXN within 4 years, and those above 4520 MXN within 6 years. However this is unlikely to be an issue for our study for two reasons: i. Conditional on default the average debt defaulted on in our sample is 15,635 MXN; ii. we can compute the default episodes for each individual and if that were to be an issue we should see a downward trend in the number of defaults per individual, this is clearly not the case in our data.

the ability to sue the client and in terms of reserve requirements). We also present results for the share of credit cards in default, defined as the ratio of the number of cards in default over the total number of active cards. Measuring default as a share of cards helps easing concerns about default being driven mechanically just by the simple fact of having more cards to default upon for those above the threshold. In the analysis, we show that all results go in the same direction. The risk measures we use in Panel B include credit cards that are active at application as well as those that were closed within 12 months before application, but not cards opened after application to Bank A. It turns out that the environment we study is risky: on average 5% of applicants had defaulted in some card before they applied for the new card. The share of cards in default is 4%. Columns 2 and 3 show that these realized risk measures are inversely related to the credit scores, as would be expected. In the last column, we report tests of equality of means across subsamples and find that these differences are statistically significant.

Finally, Panel C displays some of the variables related to the application process. Bank A's data shows that around 30% of all applications in this more restricted range were approved. It also shows that applicants request larger lines than are approved. While on average applicants requested 20,599 MXN, approved applications received on average a credit limit of 15,667 MXN (940 USD). The fact that people are applying, that they get 25% lower limits than requested, and that they accept interest rates of 37% per year (this number does not include fees, APRs are higher, not in table) may suggest that they are liquidity constrained. Note also that given that total debt is 34,746 (and the limit of revolving lines is 32,959 MXN, not shown in table), card approval represents a substantial increase in borrowing opportunities.

How do these numbers compare to those of Mexican cardholders in general? We can compare some of these statistics to those of a random sample of Mexican cardholders in June 2010 displayed in Castellanos et al. (2017). It turns out that the characteristics of our sample are similar to the characteristics of their random sample in 2010. Mean tenure in the CB is 6.5 years vs 8 in our sample, 50% are male vs 58% in our sample, people have an income of 14,300 pesos per month vs 12,910 in our sample, and the number of credit cards is 1.9 on average vs 1.7 in our sample. The sum of all credit lines is larger for Mexican cardholders however, at 53,000 pesos vs 34,314 in our sample.

3 Empirical Strategy and Methodology

The wealth of data and the clear rules for obtaining a Bank A credit card allow us to use a fuzzy regression discontinuity design, with the credit score as a running variable, to estimate the causal effect of additional credit on default in all loans, and in sequentially previous loans (Thistlethwaite and Campbell (1960), Hahn et al. (1999) and Imbens and Lemieux (2008)).[15] The identification requirements underlying this methodology are that there is a discontinuous jump at the threshold of the probability of getting the card, and that all other observed and unobserved variables are a smooth function of the running variable at this threshold. In this section, we show that in terms of observables these requirements hold in our context. We can therefore identify and estimate the Intent-to-Treat (ITT) effect by the following equation:

$$y_{it} = \alpha + \beta \mathbf{1} \left(score_{it} \geq \overline{score}_t \right) + f(score_{it}; \nu^-, \nu^+) + X'\xi + \epsilon_{it}, \tag{1}$$

where the parameter of interest β is the local, to the threshold, Intent-to-Treat (ITT) effect. This parameter is identified by the assumption that ϵ_{it}, as well as all the possible observables $X's$, are continuous at the threshold $\overline{score}_t$. Following the RD literature, we accommodate potential differences away from the discontinuity point by using a polynomial in the running variable indicated by the function $f(.)$, where we allow the shape of the polynomial (but

[15] A similar approach to ours was used in the study of mortgage securitization by Keys et al. (2010).

not the degree) to vary on the left (ν^-) and right (ν^+) of the discontinuity. For the main results we use a cubic polynomial, but provide a series of robustness checks with respect to the choice of $f(.)$. In practice, since we have two discontinuities along the credit score, we estimate 2 different $ITTs$ within a single equation, one for each threshold. The vector of controls X includes calendar month dummies as well as dummies for the number of active cards and other types of loans at the moment of application (these latter set of dummies are included when analyzing outcomes only). Somewhat liberally, throughout the paper we will refer to applicants to the left of the threshold but close to it as controls and those close but to the right as treated.

Since our design is a "fuzzy" one (i.e. not all applicants above the thresholds are given a card), in order to estimate the effect of actually obtaining a card (i.e. the local ATT) we instrument the endogenous variable (i.e. Bank A's approval of the credit card application, CR_{it}) with the indicator variable that is equal to one if the applicant's score is above the corresponding threshold. The two-stage representation of this strategy is the following:

$$CR_{it} = \alpha_1 + \beta_1 \mathbf{1}\left(score_{it} \geq \overline{score}_t\right) + f(score_{it}; \theta^-, \theta^+) + \epsilon_{it}, \tag{2}$$

$$y_{it} = \alpha_2 + \beta_2 CR_{it} + f(score_{it}; \gamma^-, \gamma^+) + \eta_{it}. \tag{3}$$

We will discuss the parameters of interest as we go along with the analysis.

3.1 Validity of the Design

We present a series of visual and formal tests of the main assumptions underlying the RD design: first, we show that the probability of obtaining a credit card is discontinuous at the thresholds; second, that the density of the credit score (the running variable) is continuous around the thresholds; and third, that an extensive set of applicants' characteristics are continuous at the thresholds.

3.1.1 Discontinuous Treatment Probability

Figure 4 shows that indeed the approval probability has a large and precise discontinuity at the thresholds. On average, the probability of obtaining a credit card to the left of the thresholds is virtually 0, while it sharply jumps to about 0.45 just to the right of the discontinuity. Such differential probability of receiving a credit card is fairly similar over the two different score thresholds. In fact, we cannot reject the hypothesis that the jumps are statistically the same in the first column of Table 3. It is also clear that our design is a fuzzy discontinuity design, in which not everyone just above the discontinuity point gets a new credit card. The fuzziness in the design, on the right hand side of the thresholds, arises from a set of extra rules imposed by Bank A: in terms of income, existing credit lines and limits. However, what is crucial for identification is that the sequence of conditions imposed starts off with the credit score. That is why all other applicants' characteristics are balanced at the thresholds as we show below.

3.1.2 Smooth Density of Applicants at the Thresholds

Another assumption that needs to hold for the RD design to be valid, is that applicants do not have the ability to *precisely* manipulate their credit score in order to *precisely* sort themselves around the discontinuity thresholds (Lee and Lemieux (2010)). We argued above why this is a reasonable assumption in our context. Figure 5 presents the empirical evidence that supports the validity of this assumption. The histograms of the standardized credit score in our pooled sample and in each subsample show that there are no noticeable discontinuities in the density at the cutoff values. A parametric McCrary (2008) test cannot reject the null hypothesis of no discontinuity, with p-values of 0.29 and 0.42 for the 670 and 700 cutoff samples, respectively.

3.1.3 Smoothness of Pre-determined Characteristics at the Threshold

A third test of the validity of the research design is that the average characteristics of the applicants on both sides of the discontinuity are statistically identical. We perform such tests on the available variables, graphically in Figures 6 and 7, and in a regression framework in Table 2. For brevity, we present the figures for the pooled sample, while the tables produces the relevant statistics for both the pooled sample and the different thresholds. The corresponding figures for each threshold separately can be found in the Appendix.

We cannot detect any statistically significant difference across the thresholds in applicants' traits or the status of their loans at the time of application. Demographic variables include: gender, income, amount requested, tenure at the credit bureau, number of credit cards 30 days before application, and total debt at the December 2010 snapshot. Note also that the economic magnitudes of the threshold coefficients are small. Perhaps more importantly, Panels E and F of Table 2 show balance in pre-determined default and delinquency measures. These variables are defined in the note of Table 1.

Overall, these results lead us to conclude that the RD methodology is valid as individuals in the neighborhood of the threshold are essentially identical. We apply this methodology in what follows.

4 Main Results

4.1 Effect on Credit Card Availability and Persistence of Treatment

Although sometimes overlooked, showing that the likelihood of getting Bank A's card is much larger to the right of the threshold is not the same thing as showing that only applicants to the right of the threshold get a card *in the market*. One might expect that rejected applicants look for loans elsewhere. In fact, one important difficulty of measuring the causal effects of credit is the widespread availability of credit, which allows the control group to access other loans and therefore to dilute the credit/no-credit comparison. An important advantage of our paper is that we are able to measure "non-compliance" for the control group using the universe of formal loans. This allows us to look at the persistence of the difference in the number of credit lines between the treatment and control groups.

Using CB data, the second column of Table 3 confirms the discontinuity in the probability of approval for the new credit card for individuals who are just above the specified threshold in terms of credit score. The probability of obtaining a new card increases by about 45pp for the pooled sample, while the number of credit cards owned mechanically increases by about 1 for those who obtain the new card. Although we do not observe the immediate increase in credit limits and debt due to the data structure, a back of the envelop calculation suggests that the ITT of limit increase is 7,200 MXN (0.45 cards×16,000 approved limit), about a 21% ITT increase of total limit and a 47% increase for those that are actually approved, a substantial increase.

Interestingly, the treatment-induced difference in the number of cards is highly persistent. Column 2 of Table 3 shows that one month after the application the differences are about the same for the two thresholds, an are still present over 12 and even 18 months after application, at about 37pp on average. We also show the monthly evolution of the number of credit cards for each cutoff in Panels (b) and (c) of Figure 8. The last two columns of Table 3 show that there is some catching-up in terms of non-card loans for rejected 700-threshold applicants. While, if anything, the treated 670-threshold applicants have more non-card loans (although not significant 18 months later).[16]

[16]Ideally, since previously existing lines could change as a result of treatment, we would also like to study the intensive margin of loan use, as in the total size of debt and credit line. Unfortunately, we face two problems in attempting to do this. First, since CB data does not contain information on debt and size of credit lines with monthly frequency, we are forced to use the June 2013 snapshot, with the consequence that a substantial amount of time elapsed since application. Second, as we described earlier, there are several reasons to believe that the quality of the variable measuring debt is lower than that of default and the number of loans. Nonetheless, with these limitations in mind, focusing on credit limit for loans active at application and using the same RD specification we estimated that our 670-threshold treatment group has 2,231 MXN larger limit in 2013 than their respective controls in loans active at application, and that the 700-threshold treatment group has 2,379 MXN lower limit.

4.2 Effect on Delinquency and Default

We now present our main results and answer three questions. First, what is the causal effect of being awarded a new credit line on default? The answer cannot be settled by theory alone. On the one hand, many models of moral hazard-driven-default, or even purely mechanical models of debt overhang and non-strategic default with income shocks, suggest that more debt leads to higher default. One the other hand, one could think of a model in which higher liquidity leads to lower default, either by facilitating more productive investments or simply by providing the ability to better smooth shocks.

The second question we address is the following: does the credit-default elasticity vary by credit score? Again, the answer is not obvious. Scores are meant to measure the level of risk, not behavioral responses. Einav et al. (2016) show that health scores do not predict individual's utilization response to kinks in the budget set. On the other hand, we show that in our context the credit score *is* predictive of default responses to credit increases.

The third, and most interesting, question is about sequential banking externalities. To what extent is sequential banking a quantitatively important phenomenon as reflected in higher default for sequentially prior banks (and credit lines)? Can the induced default be high enough to merit discussing the no-universal-default regulation?

4.2.1 Overall Effects: All Lines

Table 4 presents the effect of Bank A card approval on our measures of delinquency and default. We use equations (2) and (3) to estimate these effects. We first focus on all cards, regardless of whether they were active at the moment of application to Bank A or opened after.[17] We present results regarding cumulative delinquency and default within the first 6 and 18 months after application to give a sense of default dynamics. The dynamics of default may tell us something about the underlying causes of default. For instance, default

[17] These variables include the card from Bank A if it was awarded, or any new loan opened after application. If a loan was closed after application but before our measurement, we set the variable to zero if the loan did not default during that period and to one if it did.

may arise from people slowly accumulating debt and then finding it difficult to pay it back (either because of low repayment effort or from a negative income shock). In this case, we should expect several months to pass before default happens. Default could also arise from purely strategic reasons: once a borrower has several credit lines, she can afford to default in one of them while using the non-defaulted one to access credit and smooth consumption.

To answer to our first question: the effect of an additional credit line on default is large and robust across the different measures of default and delinquency. The answer to the second question is that the effects are completely concentrated in the 670-threshold applicants, but they take time to materialize. For the 670 group, 6 months after application, the probability of default is 5pp higher for the ITT, and 11pp in the ATT, but imprecise when the outcome variable is the share of cards in delinquency or default. They are sharper and larger 18 months after application for all measures of default. Focusing on default, column 7 shows an increase of 12pp in the probability of defaulting on any loan, out of a control baseline of 24pp. This is an ITT effect on default of 51% as a proportion of mean default. The ATT effect amounts to a doubling of the probability of default. At such a longer horizon, we also find large effects when using the share of cards in default, with a magnitude of 45% of the control average cumulative default. While the magnitudes of these effects are similar to those found in other credit markets (e.g. Dobbie and Skiba (2013) find that a $50 larger loan is associated, within 2 weeks, with a 32% decrease from the mean default rate), we find an increase in default for the lower score clients. In terms of the absolute size of the effect, we find a slightly larger percentage change in default probabilities. However, in our setting the increase in credit limit is much larger and we measure cumulative default over the next 18 months after origination.

Interestingly, the sharp increase in default is not present in the 700 sample. The point estimates are 2.5 to 10 times smaller, negative, not significant at the 1 percent confidence level (but 3 out of 8 are significant at the 5 percent level). We find that card approval leads to a 2.6pp *lower* share of cards in default. This is consistent with Dobbie and Skiba (2013) who find a negative correlation between increases in the size of payday loans and default. One possibility is that this additional credit card allows for better liquidity management, e.g. paying one card with the other credit card and vice versa to avoid delinquency while coping with temporary shocks. We explore propensities to borrow in the Appendix. Column 1 of Table B1.5 shows that borrowers just above the 700 cutoff do borrow from the new card, but about half than borrowers in the 670 group (we estimate an ATT of 1829 MXN vs 3295 MXN). Debt carried in the new card amounts to 30% and 13% of monthly income for the 670 and 700 clients, respectively.[18] The average masks the fact that the propensity to borrow for the 670 group is concentrated in the right tail of debt. The higher score clients appear to be much less likely to borrow large amounts out of their credit cards (see columns 2 and 4). All in all, the propensity to borrow out of the same increase in line is much higher for 670 applicants.[19]

These results show that heterogeneity is important. We can reject the hypothesis that the treatment response is the same for the 670 and the 700 groups in the long-run with p-values below 0.01 for all definitions of default. This heterogeneity has several implications. The first is that it cautions against the tendency to generalize conclusions from the study of a narrow population (a necessity in RD designs), but also about the dangers of pooling populations broadly. One would reach a very different conclusion depending on which segment of the population is studied. In fact, we cannot reject the null hypothesis that the effect is zero if we pool the populations from both cutoffs. This may in fact provide a clue to reconcile the different findings in the literature: looking at different populations Dobbie and Skiba (2013) find that larger loans lead to decreased default, Adams et al. (2009) find that larger loan sizes and interest rates lead to more default, while Karlan and Zinman (2009) find that larger

[18]For this calculation we took the ratio of column 1 in Table B1.5 to the administrative income recorded in 2010 from Table 1.

[19]We do not have panel data on borrowing, but we would conjecture that the speed of debt accumulation is faster for 670 applicants, suggesting that they "need" the money relatively more, while a higher fraction of the 700 applicants may use the card as precautionary credit line to smooth out temporary shocks.

interest rates are only weakly –or, depending on the definition of default, not– associated with higher default. The second substantive implication of the documented heterogeneity is about financial inclusion. Going after lower credit score applicants is much riskier. Even if the difference across cutoffs is only 30 nominal score-points, the elasticity of delinquency to credit seems to be steeply decreasing in that segment. A third point is that the credit score is a good predictor not only of the likelihood of default in the next 12 months –for which it was built– but of behavioral responses as well, as reflected in the credit-default elasticities.

Higher default is not the same as lower profits. One may wonder if the 670 clients are unprofitable for the bank to go after. Unfortunately, we do not observe measures of profitability, but two facts point towards low profitability. On the one hand, a simple back of the envelope calculation shows that the bank needs about 1400 extra MXN of average daily balance uniformly during 18 months for the 670 group to pay for its higher default through interest payments.[20] This does not happen in practice. Perhaps more persuasive is the fact that, after this experiment with lower 670 credit score cutoffs, Bank A increased the cutoff, strongly indicating it was not profitable for them (they confirmed with us that this was indeed the reason).

We explored two other sources of heterogeneity: how credit-default elasticities vary with predetermined average credit line utilization and with the number of pre-existing active cards at application. Median credit card line utilization is 53 percent in our sample in the 2010 snapshot. One may conjecture that both, the liquidity management and the debt burden/moral hazard problems, are larger for those more heavily leveraged. Table B1.2 and Table B1.3 estimate the same specifications as above but using the top 30% sample of applicants with highest utilization in December 2010 (those above 76% of total line utilization). We find effects that are twice as large for the 670 applicants, and more accentuated nega-

[20]Conditional on credit card default, the average amount defaulted in the 2010 snapshot is 7,770 MXN. Given that treatment induces close to 10pp more default in the 670 group 18 months after application and that the yearly interest rate in Bank A's cards is 37%, it turns out that 1400 MXN $\times$ 18 months $\times$ (0.37/12) monthly interest = 7,770 MXN $\times$ 10% default.

tive effects for the 700 applicants. This latter results, corresponding to the 700 applicants, lends some support to our conjecture that the negative effect on default for these higher score clients could be due to better ability to smooth out temporary shocks, i.e. highly leveraged clients would not be able to repay without the extra credit line from Bank A. We also estimate higher credit-default elasticities for applicants with lower number of cards at application; these elasticities are twice as big for those that had zero cards at application than for those holding 2 credit cards for instance (see Figure B1.5 in the Appendix).[21]

4.2.2 Externality Effects: Pre-existing Lines

To what extent does approval by Bank A triggers default on pre-existing loans? Using our RDD design, we estimate no statistical difference between our control and treatment groups in the size of pre-existing credit lines (i.e. those already active at the application date). This means that if we see any default effect on those pre-existing credit lines it will entirely be due to the existence of an externality and not to the direct effect of an increase in credit limit on those lines.

Table 5 presents our main results for externalities, while the right column of Figures 9 and 10 do it graphically. We use the same specification and definitions as in Table 4, but the estimation is performed on the sample of cards that were active at the moment of the application. The rightmost columns look at spillovers to other types of loans apart from credit cards, including mortgage loans, personal loans, auto loans, and all others. As found before, there is substantial heterogeneity across the credit score distribution. For the 670 group, we find large and statistically significant effects on default on pre-existing cards and pre-existing non-card loans for all measures of default/delinquency. For instance, the probability of default in cards increases by 8.7pp from a control comparison of 20pp. This

[21] One rationalization of these results is that there exist different types of borrowers which differ in their effort cost of repayment. Consumers with higher cost of repayment will choose lower effort –i.e., will default more often– which in turn implies both: lower holdings of credit cards on average since banks are more likely to reject card applications of applicants with bad credit histories, and that those with lower credit scores have higher credit default elasticities. That is: low scores and low holdings of credit cards are a signal of high cost of effort types.

is an ITT increase of 43%, and an ATT of 92% for credit cards. For non-card loans the analogous effects are smaller but still statistically and economically significant, at 22% ITT and 48% ATT for loans active at application (columns 7 and 8). For the 700 group, we find once again that the extra card lowers default, but the effects are small and less statistically robust across delinquency/default definitions.

How can we interpret these effects? First, we find these effects to be large enough to significantly hinder the profitability of loans made by previous lenders. Just for comparison purposes, during the last recession in the US delinquency rates increased from 3.7% to 6.5% for credit card loans, from 2.8% to 4.7% for consumer loans and from 1.6% to 10.9% for residential mortgages.[22] One can also benchmark default by referring to the official stress testing of the Mexican Financial Stability Committee http://www.cesf.gob.mx/en/CESF/Publicaciones_e_informes. In its 2011 report they simulate extreme adverse scenarios in default probabilities for the different loan segments and estimate how this would impact bank solvency. The first noteworthy fact is that their extreme "adverse scenario" corresponds closely to our estimated default externalities: default probabilities for revolving and non-revolving consumption loans would increase by 78% and 56% respectively, so regulators think these are big effects.[23] The second noteworthy fact are the implications of such default increases for bank's health: "Banks estimate that... such an increase would decrease loan growth rates by 40% relative to the base realistic scenario" and "would increase expected losses, under the assumptions, in such a way that the capitalization index of the banking sector estimated in December 2012 would decrease from a projected 15.8% in the base scenario to 12.4% in the adverse scenario" (the regulatory minimum is 8%). Under this light our estimated default magnitudes could be really meaningful for the stability and growth of the financial sector.

[22]These figures were obtained from series DRCCLACBS, DRCLACBS and DRSFRMACBS provided by https://fred.stlouisfed.org.

[23]These correspond to the actual increases in Mexico during the financial crisis: 7.6pp and 6.4pp respectively — in percentage points — from December 2006 to December 2008.

It is very unlikely that first lenders find these magnitudes of default convenient and, say, recover the losses from default with extra late payment fees or overdraft fees. However, since we do not observe fees we cannot test this formally. Increases of 7pp-9pp in default will likely bankrupt most banks. Our estimates highlight the importance of sequential banking for the stability of the banking system. Sequential banking is a potentially negative aspect of competition for the same set of clients. In a similar vein, Petersen and Rajan (1995) find that in the US more banking competition is associated with less financing to small firms, presumably because having more banks makes loans less profitable. Although in our case this arises through a sequential banking mechanism, they do not identify the exact channel. Competition of course has many positive aspects. Second, one may ask why banks have not introduced covenants that restrict sequential banking behavior. We mentioned above that in Mexico such covenants are illegal, and that in the US they have been restricted by regulation. Furthermore, in our interviews with Mexican banks we discovered that while they are aware that sequential banking is an issue, for lack of a good empirical design they do not know how large these externalities are. Third, the presence of externalities does not imply that banks lose money. In fact, in the subgame perfect equilibrium of Bizer and DeMarzo (1992) banks anticipate that borrowers will bank sequentially, and since they have the option of not lending, they do not lose money in equilibrium. The existence of such externality does not imply negative profits in equilibrium, but it still causes inefficiency, as in Bizer and DeMarzo (1992) where equilibrium interest rates are larger, borrowers are more indebted, there is more default, and less efficiency than there would be if there was only one lender. Bisin and Guaitoli (2004) find that multiple banking could in fact help sustain high profits for incumbent banks.

One may wonder why the estimated default magnitudes are large. A possible explanation is that the treatment itself is large. The credit limit awarded is larger than the consumers' monthly income. For comparison, Dobbie and Skiba (2013) operate with increases of 50

USD while we estimate the effect of increases of 940 USD. One way to interpret the effect by the amount of credit limit awarded is to estimate equation 3 using credit limit as an explanatory variable instead of CR_{it}. Table B1.4 in the Appendix does exactly this. We find that an extra 1,000 MXN (60 USD) is associated with an ATT effect on the probability of default of 1.5pp for the 670 group. Another explanation is loan targeting. This focuses our attention precisely on the heterogeneity of the effect: all the positive effect we found is concentrated on "extra-marginal" applicants that Bank A previously identified as not credit worthy. In contrast, Karlan and Zinman (2009), Adams et al. (2009), and Dobbie and Skiba (2013) focused on populations that were typical and even repeat clients of the lender.

Finally, one may want to know if borrowers default on all their loans or just on a subset of them, and in the latter case which ones they choose to default on. Some models, as in Parlour and Rajan (2001), assume that the larger part of the cost of default is a fixed cost (e.g. lost reputation) and does not vary with how many loans one defaults upon. Given this assumption, they predict that people will default on all their loans or in none of them. This does not bear out in our data. Conditional on defaulting on at least one of them, borrowers that have 2 or more active loans typically default only in a strict subset of them. In fact only 2% of borrowers default on all their loans in a span of 3 months after the first loan defaulted. So, default *is* selective. This raises the question "which loans are borrowers more likely to default on?". Table B1.6 in the Appendix presents evidence consistent with borrowers being 10 percent more likely to default in their oldest loans, 20 percent less likely to default on collateralized loans, and neither more nor less likely to default in larger loans (these are just correlations and we don't want to stress it much).

5 Robustness

We consider two robustness tests to our results. The first is with respect to the specification of our main estimating equations. Figure 11 shows that our results survive when using quadratic or cubic polynomials to approximate the function $f(\cdot)$ in equation (3), restricting the sample to applicants with scores 15 points above or below of their respective cutoff, or using local linear regression with Imbens and Kalyanaraman (2011) bandwidths. The second set of robustness tests address the concern that different macro shocks may have happened when the 670 or the 700 threshold cutoffs were operational. We perform a battery of robustness tests to rule this out. Before we go to these tests note that, regardless of this potential confound, we have a valid treatment-control comparison within threshold group, since macro shocks will affect applicants across threshold in the same way as they applied at the same time. Recall also that our regressions include calendar month dummies, but the fact that results almost do not change if we exclude them suggests that the interaction of the treatment with seasonality is small, if any.

The first test exploits the fact that the periods of time of the different threshold regimes are not far apart: the 700 group applied from January 2010 to April 2011, while the 670 group applied immediately after that, from June 2011 to November 2011. This by itself casts doubt on the importance of differential time trends within such a short period of time. Nonetheless, Figure 12 performs the estimations in a sample restricted to close months. That is, it compares our main ATT results (hollow dots) against estimates obtained using the 3-month (or 2-month) sample, which uses data from applications made in February, March and April 2011 (or March, April 2011, respectively) for the estimates at the 700 threshold, while it uses data from June, July and August 2011 (or June and July 2011) for the 670 estimates. Although this reduces our sample size by two thirds (or three quarters) of the original size, the estimated effects are similar to the ones we presented in the baseline estimations.

A second piece of evidence comes from applicants' characteristics during our sample period. One may also worry that the selection of applicants during the 700-threshold period is different from that during the 670-threshold period. This is a priori unlikely, as the threshold change was internal to the risk department and neither the applicants nor the branch managers knew about it. Figure B1.8 in the Appendix plots the time trends of averages of (a) the score of the applicants, (b) the self-reported income of applicants, (c) age, (d) gender of applicants at the different calendar months. To have the same units, each variable is normalized so that January 2010 applicants have a mean of 100 in each variable. The vertical line indicates when the 670 period started. There are no pronounced trends in any of these variables, indicating that the selection of applicants is similar across time.

Figure B1.9 in the Appendix presents our third test. It compares cumulative default time trends for applicants who are "always-controls" regardless of the threshold regime, i.e. those with a score in the [640,660] range. We do this in order not to confound a differential treatment effect with a differential time trend. If default time trends during the 700 threshold regime period were different from default time trends during the 670 threshold regime, this would show up in different default levels for the always-control group at these different periods.[24] Figure B1.9 shows that this is not the case. It presents the regression-estimated difference in cumulative default between applicants with scores in the [640,660] range in the 700 regime and applicants with scores in the [640,660] range in the 670 regime. We find no difference in cumulative default rates for the always control group in the two different threshold regimes.

Overall we find that our results are robust and our identification strategy within and across thresholds regimes to be valid.

6 Conclusions

This paper makes several contributions to our understanding of credit markets. First, using country wide data from Mexico on the universe of loans, it documents that sequential banking is prevalent. Second, using exogenous shocks to the credit availability of borrowers and data from a large bank with presence in all 32 states, it estimates credit-default elasticities across

[24] This exercise assumes that the causes of default that may affect people with scores close to the 670 and 700 thresholds are similar to the ones affecting applicants with scores between 640 and 660.

all types of loans. In this margin, this paper generalizes the important papers by Karlan and Zinman (2009), Adams et al. (2009), Dobbie and Skiba (2013), and Agarwal et al. (2017) by focusing on a low-inclusion country, all types of loans, and a lender with large geographical breadth serving costumers spanning a large income range. Thirdly, Bank A's pursuit of extra-marginal borrowers gives us a scarce and valuable window into the efforts and difficulties of the much discussed "financial inclusion". We document that default elastiticies increase fast as we go down the credit score; this may help explain why there seems to be a substantial amount of rationing in the financial market just below certain levels of credit worthiness (e.g. Adams et al. (2009)). In particular, from our calculations and Bank A's revealed behavior it appears that sliding down the score distribution was not profitable.

Fourth and connected to the above, although we use an RD design, the fortune of having different cutoffs allows us to estimate more than one local effect. This enables us to document for the same country and even the same bank that external validity issues could go a long way to reconciling the differences across studies of moral hazard in the credit market. Finally, we provide the first causal quantification of sequential banking externalities and find that they are really substantial, leading to almost a doubling of default. This has direct policy relevance for the design of no-universal-default clauses and for the efficiency of financial markets.

Figure 1: Elapsed Time Between Loans

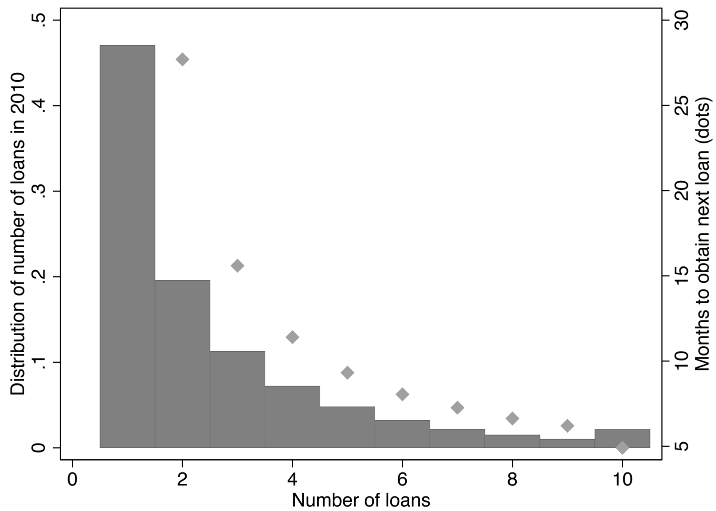

Notes: This Figure plots a histogram of the number of active loans held by a random sample obtained from the universe of card holders in Mexico in 2010, as well as the average number of months between the opening dates of the 2nd and the 1st loans, the 3rd and the 2nd loans, etc. The prevalence of sequential banking manifests in the fact that starting from the 3rd loan onwards it takes a few months from sequentially getting an additional loan. Confidence intervals are not reported since they are small enough to get confounded with the dots.

Figure 2: Timeline

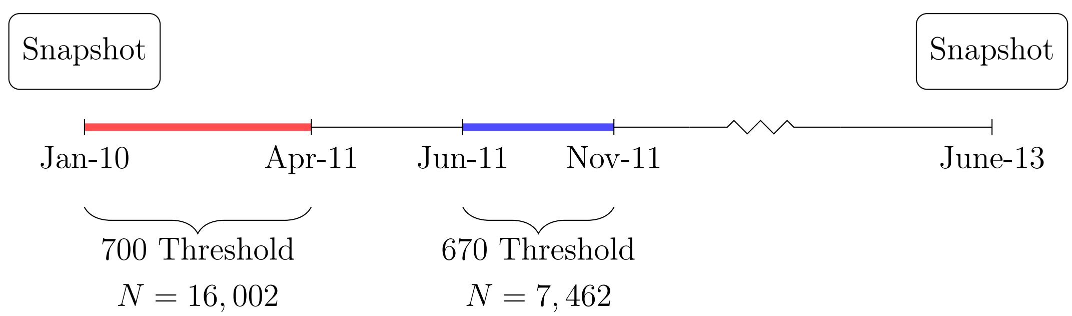

Notes: This figure represents a timeline of Bank A's application process and shows the dates when different thresholds were used in the approval process. It also includes the dates of the different snapshots we received from the Credit Bureau.

Figure 3: Credit Score vs Default

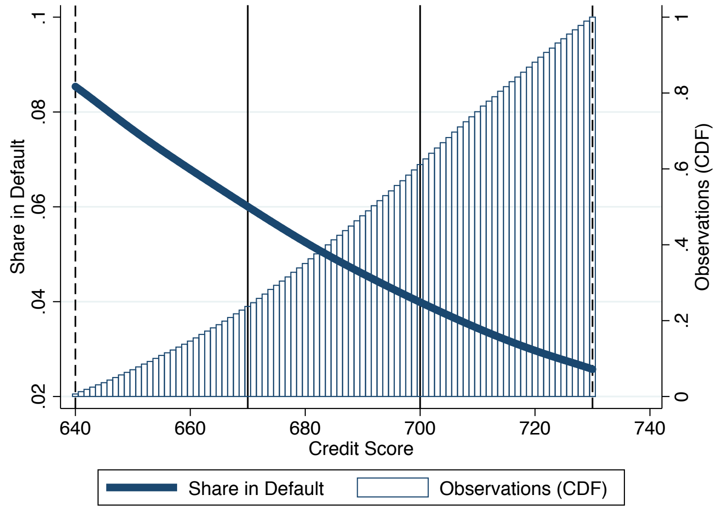

Notes: The figure shows the relationship between credit score (x-axis) and default probability (left y-axis) using data from the entire sample of applicants to Bank A's credit card. For each value of the credit score, we compute the share of loans that were ever in default during the 12 months after the date of application (also when the credit score was measured). The vertical right axis shows the cumulative distribution of loans in the 640-730 range of the credit score.

Table 1: Summary Statistics

	All	Score Cutoff 670	700	670 = 700
Panel A: Demographics				
Income (MXN)	12910	11055	14199	0.000
	(11162)	(10033)	(12016)	
Male	0.58	0.56	0.58	0.320
	(0.49)	(0.5)	(0.49)	
Panel B: Pre-treatment Credit Characteristics				
Tenure in Bureau (Years)	7.7	8.0	7.6	0.005
	(4.67)	(5.08)	(4.5)	
# of non-Bank A CC 30 days before	1.70	1.56	1.65	0.157
	(1.86)	(1.72)	(1.98)	
# of Active Credits 30 days before	3.70	3.95	3.56	0.000
	(2.88)	(2.92)	(2.97)	
Total Debt (MXN)	34746	30043	37038	0.002
	(58714)	(54697)	(61097)	
Total Limit (MXN)	34314	28924	36105	0.004
	(64620)	(58319)	(69377)	
# CC in Default†	0.06	0.09	0.05	0.000
	(0.28)	(0.36)	(0.24)	
Probability of CC in Default†	0.05	0.07	0.04	0.000
	(0.22)	(0.26)	(0.2)	
Share of CC in Default†	0.04	0.06	0.04	0.001
	(0.18)	(0.21)	(0.17)	
# CC in 2 Months Delinquency†	0.08	0.13	0.07	0.000
	(0.33)	(0.41)	(0.29)	
Probability of CC in 2 Months Delinquency†	0.07	0.11	0.06	0.000
	(0.25)	(0.32)	(0.23)	
Share of CC in 2 Months Delinquency†	0.06	0.10	0.05	0.000
	(0.2)	(0.26)	(0.19)	
Panel C: Applications				
Approved	0.33	0.33	0.29	0.003
	(0.47)	(0.47)	(0.45)	
Amount Requested (MXN)	20599	16196	22086	0.000
	(17926)	(17381)	(17533)	
Approved Amount (MXN)**	15667	16483	14698	0.014
	(12292)	(11594)	(12383)	
N	23464	1228	3229	

Notes: This table presents summary statistics of our sample of applicants. The first column reports summary statistics for the 700 and 670 sample pooled together, including applicants that had a credit score within the +/− 30 points range around the threshold at the moment of application. The next two columns report summary statistics for each of the two sub-samples (focusing on applicants that had a credit score within the +/− 5 points range around the threshold at the moment of application). Finally, the last column reports the p-value of the test of the null hypothesis that the means in the 700 and 670 subsamples are equal. Observations with a cutoff of 700 points correspond to applications made between January of 2010 and April of 2011, while observations with a cutoff of 670 correspond to applications made between June and November of 2011. Income was obtained from administrative social security data (with a match rate of 21%). Panel B presents measures of delinquency and default defined at the applicant level. Probability of delinquency is equal to one if the person has ever had any credit card with 60 to 90 days past due from the earliest month with available information of the card to the date of application. The probability of default is analogously defined, but defining default as delays in payment longer than 90 days. The share of delinquent credit cards is defined as the ratio of the number of cards with 2 months delinquency over the total number of cards. The share of credit cards in default is analogously defined. The † signs indicates that the variable was constructed using all the available past information of the credit cards with opening dates earlier than the application date, ,i.e. for previously-existing card. ** Average credit limit approved by Bank A, conditional on approval.

Figure 4: Percentage of Approved Applications by Score and Cutoff

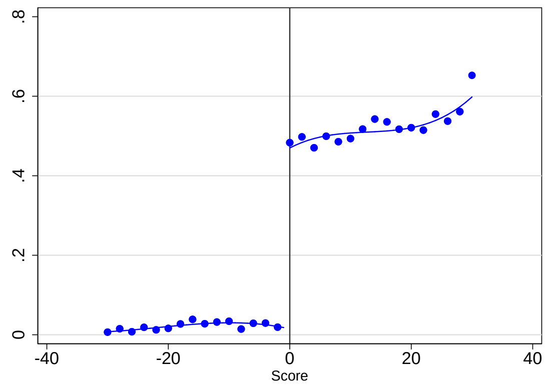

(a) Pooled Sample

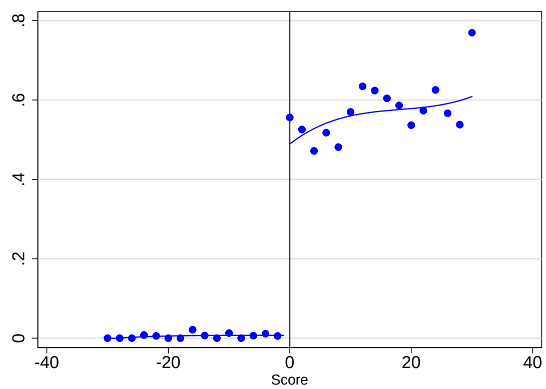

(b) 670 Sample

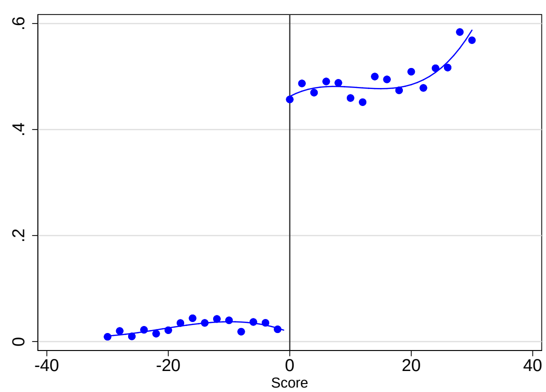

(c) 700 Sample

Notes: The figure presents the percentage of credit card applications that were approved by the bank for each pair of values of the standardize credit score between standardized scores of -30 and 30. It also presents a polynomial fit of degree three to the raw data, allowing the intercept and the coefficients of the polynomial to differ on both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The first figure corresponds to the pooled sample, while the second and third figures refer to the 670 and 700 samples, respectively.

Figure 5: Distribution of Credit Score with a Polynomial Fit

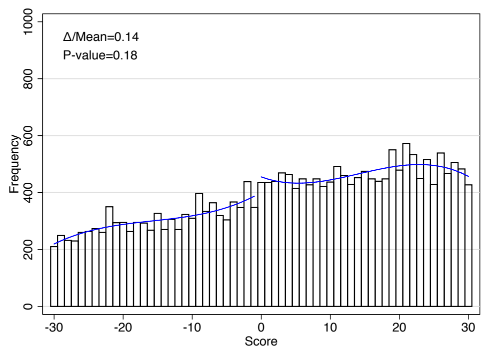

(a) Pooled Sample

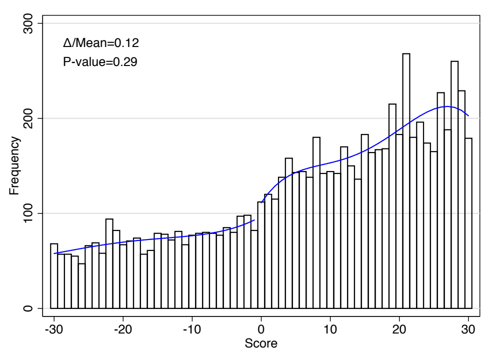

(b) 670 Sample

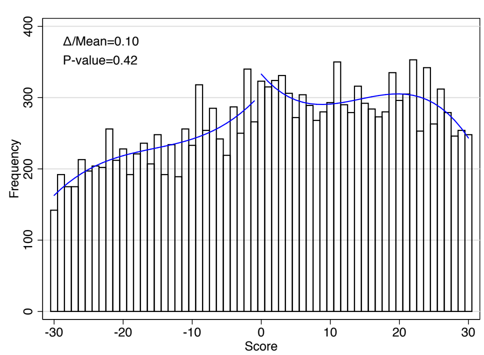

(c) 700 Sample

Notes: The figure presents the frequency distribution of credit scores in the population of applicants. The size of each bin corresponds to one point of the credit score. Panel (a) shows the histogram for the pooled sample with the score standardized so that 0 equals the threshold score for each subsample. Panels (b) and (c) show the histogram of the raw score for each subsample separately. The blue lines represent two approximating third order polynomials at each side of the threshold (for the 670 sample we included a fourth order term). We also report the value of the discontinuity at the threshold as a percentage of the mean frequency, and the p-value of the test of the null hypothesis that there is no discontinuity at 0.

Figure 6: Pre-Treatment Characteristics – Pooled Sample

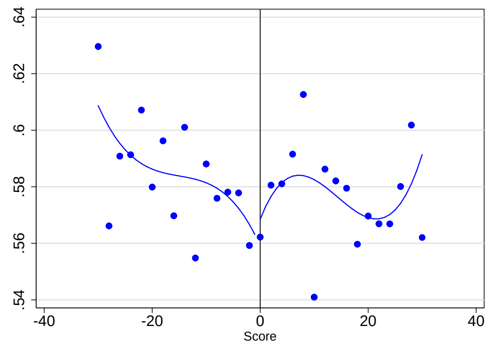

(a) % Male

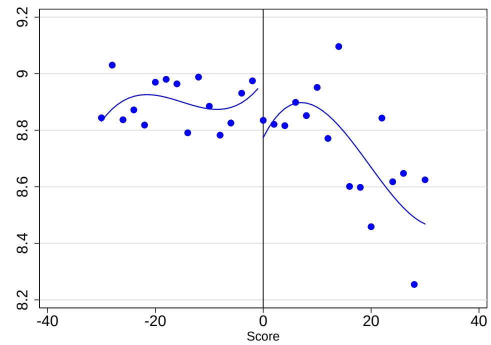

(b) Amount requested (Log)

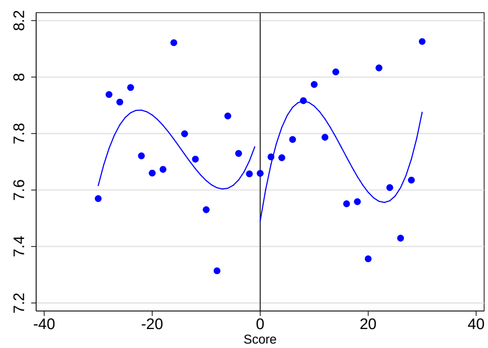

(c) Tenure in Bureau (years)

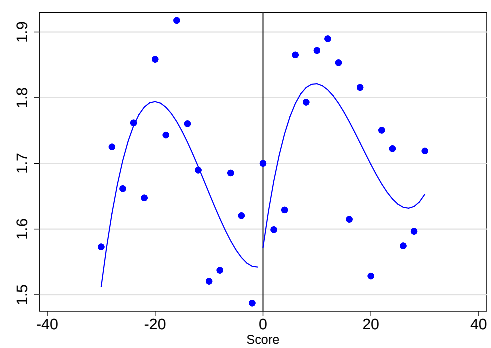

(d) # CC 30 days before

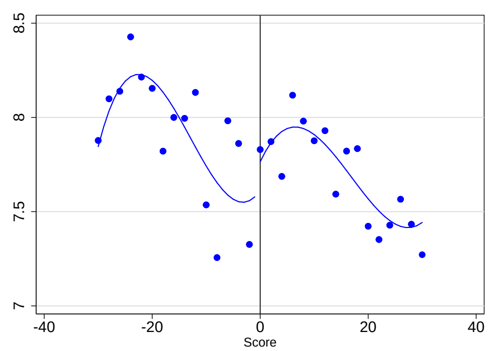

(e) Total debt (Log)

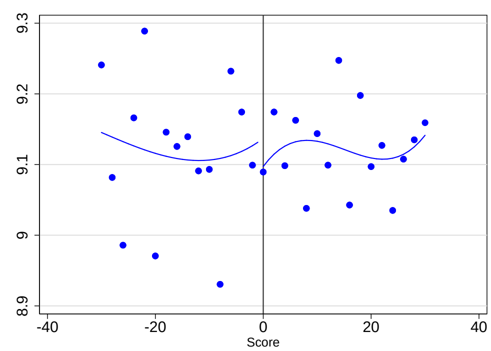

(f) Income (Log)

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. Panel (a) refers to the percentage of males in each score bin, Panel (b) to the credit limit requested at the application in logs, Panel (c) to the years each person has been in the Credit Bureau, Panel (d) to the number of active credit cards applicants had 30 days before the application, Panel (e) to total Debt in 2010 and panel in logs and Panel (f) to the applicant's administrative income in logs. Appendix B shows similar graphs split by cutoff.

Figure 7: Pre-Approval Outcome Variables – Pooled Sample

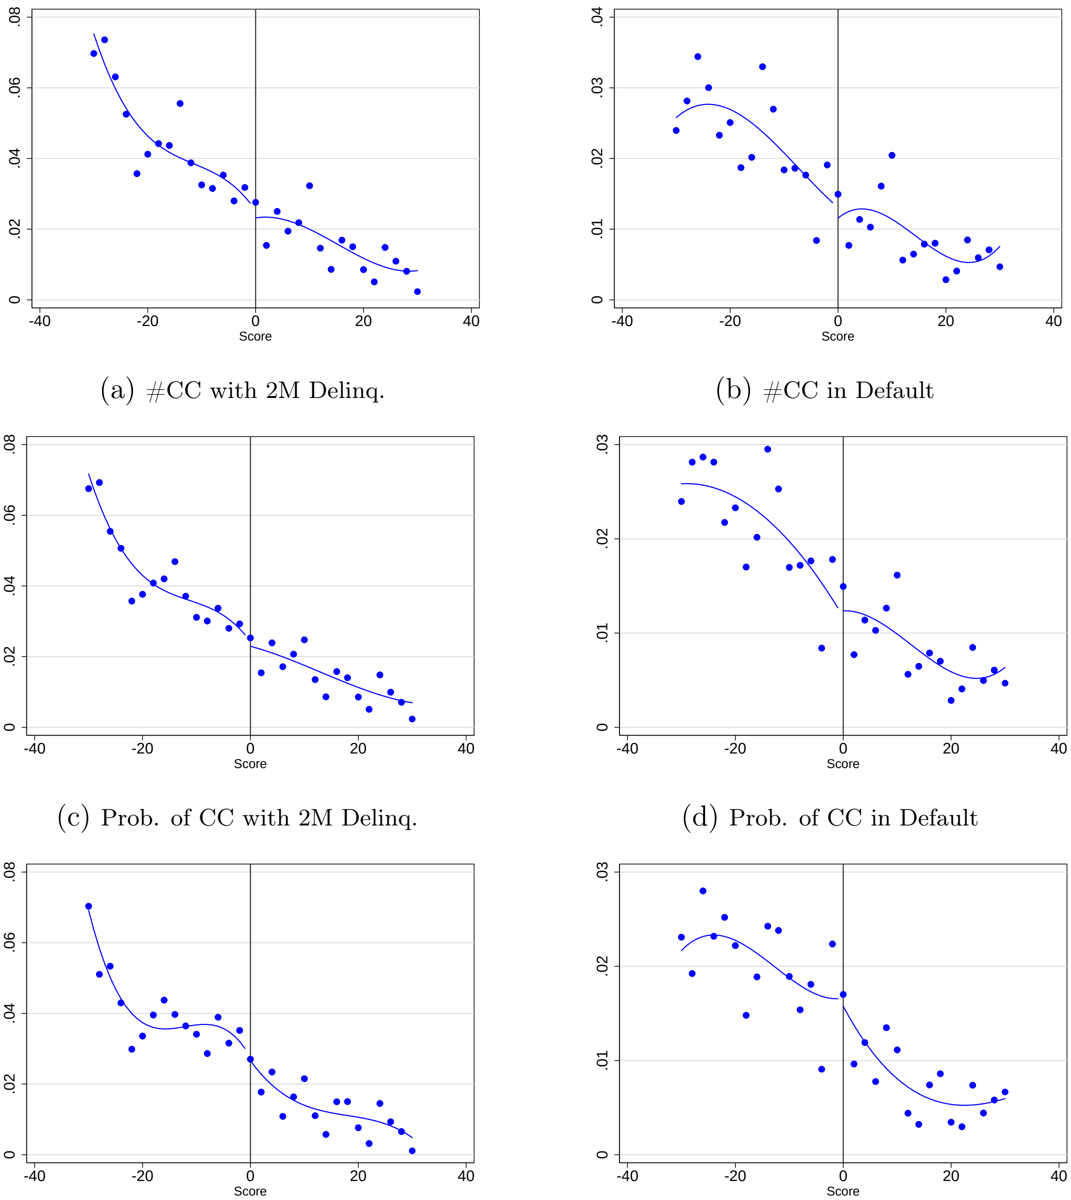

(a) #CC with 2M Delinq.

(b) #CC in Default

(c) Prob. of CC with 2M Delinq.

(d) Prob. of CC in Default

(e) Share of CC with 2M Delinq.

(f) Share of CC in Default

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. Panel (a) refers to the number of credit cards that had at least one delinquency episode from the earliest month with available information of the card to the application date, Panel (c) to the probability of a delinquency, which is defined as an indicator variable that is equal to one if the applicant has had any credit card with 60 to 90 days past due from the earliest month with available information of the card to the application date, and Panel (e) to ratio of the number of cards in delinquency over the total number of cards. Panels (b), (d) and (f) are analogous but focus on default, which is defined as late payments of 90 days or more. These variables were constructed including only credit cards that were active at the date of application. Appendix B shows similar graphs split by cutoff.

Table 2: Tests of Quasi-Random Assignment of Pre-Determined Characteristics

	Male	Tenure (Years)	$\#CC$ 30 Days Before	Total Debt (Log)	Administrative Income (Log)	Amount Requested (Log)
			Panel A: Pooled Sample			
Above Cutoff	0.009	-0.322	0.042	-0.128	-0.043	-0.100
	(0.017)	(0.249)	(0.199)	(0.138)	(0.086)	(0.097)
Mean Dep. Var.	0.58	7.72	1.61	6.92	9.16	8.93
N	23492	23492	23492	23492	4935	23492
			Panel B: By Cutoff			
Above Cutoff 670	-0.010	-0.173	0.235	0.112	-0.011	-0.093
	(0.036)	(0.473)	(0.205)	(0.152)	(0.182)	(0.257)
Above Cutoff 700	0.017	-0.423*	-0.029	-0.208	-0.041	-0.092
	(0.023)	(0.245)	(0.219)	(0.161)	(0.142)	(0.104)
			Panel C: Means[-5;-1] from threshold			
670	0.58	8.18	1.45	6.30	9.02	7.34
700	0.58	7.57	1.66	7.13	9.21	9.45
			Panel D: Joint Testing (p-values)			
670 = 700	0.578	0.586	0.147	0.063	0.918	0.998
	$\#CC$ with 2M Delinq.	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	$\#CC$ in Default	Prob. of CC in Default	Share of CC in Default
			Panel E: Pooled Sample			
Above Cutoff	-0.003	-0.002	-0.001	-0.002	0.000	-0.000
	(0.006)	(0.004)	(0.005)	(0.008)	(0.007)	(0.006)
Mean Dep. Var.	0.03	0.03	0.02	0.01	0.01	0.01
			Panel F: By Cutoff			
Above Cutoff 670	-0.002	-0.006	-0.009	0.012	0.013	0.007
	(0.020)	(0.019)	(0.023)	(0.021)	(0.020)	(0.019)
Above Cutoff 700	-0.004	-0.002	0.001	-0.007	-0.004	-0.003
	(0.007)	(0.006)	(0.005)	(0.006)	(0.005)	(0.004)
			Panel G: Means [-5;-1] from thresholds			
670	0.06	0.06	0.05	0.02	0.02	0.02
700	0.02	0.02	0.01	0.01	0.01	0.01
			Panel H: Joint Testing (p-values)			
670 = 700	0.940	0.886	0.691	0.363	0.366	0.576

Notes: This table presents the results of tests of quasi-random assignment of credit cards 30 points around the cutoff. The estimates were obtained by OLS regressions of the applicant's characteristic on a third order polynomial, allowing the intercept and the coefficients of the polynomial to differ at both sides of the cutoff. Clustered standard errors at the credit score level are reported in parenthesis. We control for cyclical and seasonal variation by including indicator variables for each month during the application period. For those observations in which an applicant did not have a credit before the application (and therefore measured delinquency is not defined), we set the variables to zero and flagged those observations with an indicator variable. In the first part of the table, Male is a dummy variable for Male applicants. Tenure is the number of years of tenure in the Credit Bureau. Number of credit cards 30 days before are the number of active credit cards the applicant had 30 days before the application date. Total debt is the logarithm of the total debt in all active credits in January 2010. Income is the applicant's income as reported to the social security. Amount requested measures the logarithm of the requested line in the application. Panels E, F, G and H use the same pre-treatment outcome variables as Figure 7. Panels A and E show results for the pooled data, Panel B and F present the results for each cutoff sample separately. Panels C and G display the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panels D and H present the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure 8: Credit Expansion and Persistence over Time

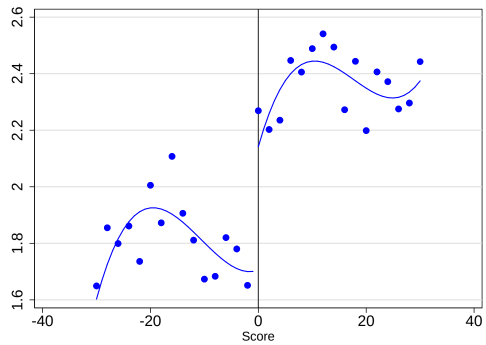

(a) # CC one Month after Application

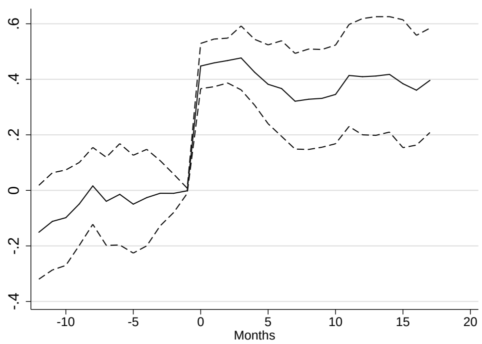

(b) CC Gap over Time for 670 Cutoff

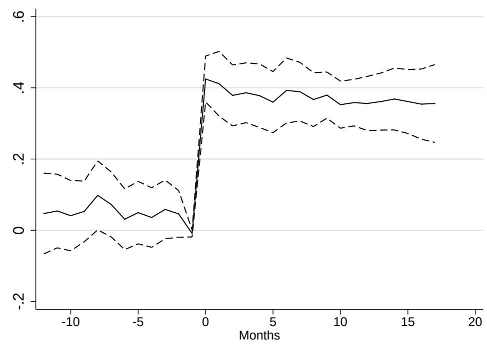

(c) CC Gap over Time for 700 Cutoff

Notes: This figure shows that treatment led to an immediate difference in the number of cards between treatment and control applicants. Panel (a) shows the number of credit cards one month after application, for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also presents a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. Panel (a) shows that applicants to the right of the threshold had about 0.4 cards more one month after application, which corroborates the fact that Bank A's approval information is consistent with the CB data. Panels (b) and (c) present OLS estimates of our main RD specification (Eq. 1) with the number of active credit cards by month as the dependent variable. The equation is estimated for each month separately, starting from 12 months before and extending up to 18 months after application. Both figures show the β coefficients and their 90% confidence intervals for the corresponding month.

Table 3: RD Estimates of the Effect of Approval on the Number of Cards

	Probability of Approval	*#CC* 1 Month After	*#CC* 6 Months After	*#CC* 12 Months After	*#CC* 18 Months After	# Credit Lines 6 Months After Excl. CC	# Credit Lines 18 Months After Excl. CC
				Panel A: OLS			
Pooled cutoffs	0.451***	0.431***	0.363***	0.371***	0.365***	0.006	-0.092
	(0.019)	(0.038)	(0.041)	(0.030)	(0.043)	(0.054)	(0.067)
Above cutoff 670	0.472***	0.448***	0.382***	0.414***	0.396***	0.234**	0.125
	(0.058)	(0.049)	(0.086)	(0.111)	(0.114)	(0.098)	(0.157)
Above cutoff 700	0.444***	0.425***	0.360***	0.358***	0.356***	-0.070	-0.166**
	(0.015)	(0.039)	(0.052)	(0.040)	(0.066)	(0.063)	(0.078)
				Panel B: IV			
Pooled cutoffs	-	0.955***	0.805***	0.823***	0.810***	0.014	-0.203
	(-)	(0.055)	(0.072)	(0.068)	(0.097)	(0.119)	(0.149)
Approved 670	-	0.950***	0.810***	0.878***	0.841***	0.497**	0.269
	(-)	(0.048)	(0.181)	(0.247)	(0.262)	(0.228)	(0.331)
Approved 700	-	0.958***	0.812***	0.810***	0.805***	-0.158	-0.373**
	(-)	(0.081)	(0.119)	(0.094)	(0.154)	(0.141)	(0.176)
				Panel C: Means [-5;-1] from cutoff			
Pooled cutoffs	0.03	1.763	1.894	1.938	1.939	2.304	2.477
670	0.01	1.511	1.548	1.575	1.575	2.658	2.412
700	0.03	1.844	2.006	2.056	2.057	2.189	2.499
N	23492	23492	23492	23492	23492	23492	23492
				Panel D: Joint Testing (p-values)			
OLS 670 = 700	0.648	0.605	0.837	0.684	0.799	0.010	0.109
IV 670 = 1	-	0.296	0.294	0.621	0.544		
IV 700 = 1	-	0.606	0.115	0.042	0.205		

Notes: This table reports the first stage estimates and the RD estimates of the effect of eligibility on the number of loans at different horizons. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. In column (1) the dependent variable is a binary variable that indicates the approval of a credit card application. In the next four columns, the dependent variable is the number of active credits cards 1, 6, 12 and 18 months after the application. In the last two columns, the dependent variable is the number of active non-credit card loans 6 and 18 months after the application. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 4: The Effect of Approval on 6 month and 18 month Delinquency and Default

	Short run (6 Months)				Long run (18 Months)			
	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	Prob. of CC in Default	Share of CC in Default	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	Prob. of CC in Default	Share of CC in Default
Panel A: OLS								
Pooled cutoffs	-0.015	-0.014	-0.013	-0.010	0.025	-0.003	0.014	0.000
	(0.020)	(0.016)	(0.015)	(0.013)	(0.021)	(0.014)	(0.021)	(0.013)
Above cutoff 670	0.053***	0.019	0.052***	0.022	0.127***	0.084**	0.122***	0.077**
	(0.019)	(0.018)	(0.019)	(0.017)	(0.047)	(0.039)	(0.044)	(0.033)
Above cutoff 700	-0.038	-0.025	-0.036**	-0.020	-0.011	-0.032**	-0.024	-0.026**
	(0.027)	(0.019)	(0.017)	(0.014)	(0.021)	(0.013)	(0.018)	(0.011)
Panel B: IV								
Pooled cutoffs	-0.034	-0.032	-0.029	-0.022	0.056	-0.006	0.031	0.000
	(0.045)	(0.036)	(0.033)	(0.029)	(0.046)	(0.031)	(0.045)	(0.029)
Approved 670	0.113**	0.040	0.111***	0.046	0.270***	0.181**	0.260***	0.166**
	(0.044)	(0.039)	(0.038)	(0.035)	(0.097)	(0.082)	(0.092)	(0.067)
Approved 700	-0.086	-0.057	-0.082**	-0.046	-0.024	-0.070**	-0.053	-0.056**
	(0.060)	(0.044)	(0.038)	(0.032)	(0.048)	(0.029)	(0.041)	(0.025)
Panel C: Means [-5;-1] from cutoff								
Pooled cutoffs	0.105	0.067	0.075	0.048	0.234	0.159	0.203	0.140
670	0.133	0.093	0.093	0.063	0.283	0.201	0.238	0.173
700	0.096	0.059	0.069	0.043	0.218	0.146	0.192	0.130
N	23492	23492	23492	23492	23492	23492	23492	23492
Panel D: Joint Testing (p-values)								
670 = 700	0.010	0.074	0.000	0.018	0.006	0.006	0.001	0.003

Notes: This table reports the RD estimates on different measures of cumulative delinquency and default during the first 6 and 18 months after the application. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. The dependent variables were constructed using information from all credit cards that were active at application as well as those opened afterwards. Probability of CC with 2M Delinquency is an indicator variable that is equal to one if the applicant has had at least one delinquency episode from the date of application to the first six months after application. Share of CC with 2M Delinquency is the share of cards that were in such a situation during the same period of time. The next two columns are analogous to the former but focus on default, which is defined as late payments of 90 days or more. The last four columns are equally defined as the first 4 columns, but consider the period of time from the date of application to the first eighteen months after application. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure 9: Effect on Long-Run Delinquency (18 Months) – 670 Sample

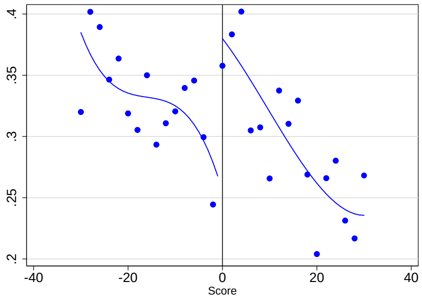

(a) Prob. of CC with 2M Delinq. †

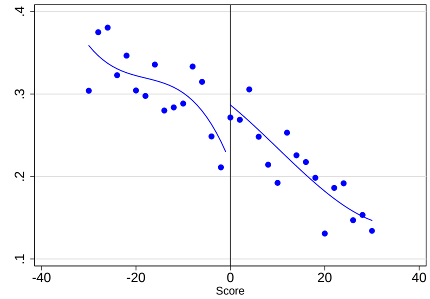

(b) Prob. of CC with 2M Delinq. ‡

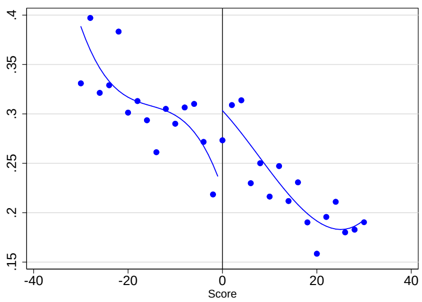

(c) Share of CC with 2M Delinq. †

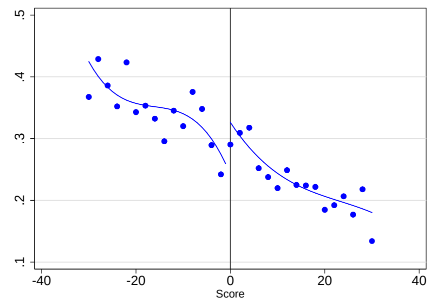

(d) Share of CC with 2M Delinq. ‡

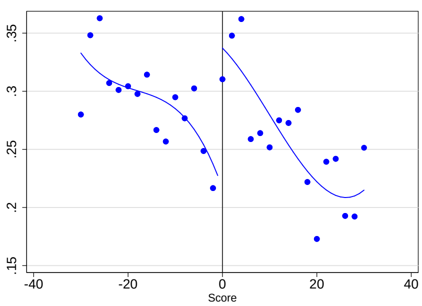

(e) Prob. of CC in Default †

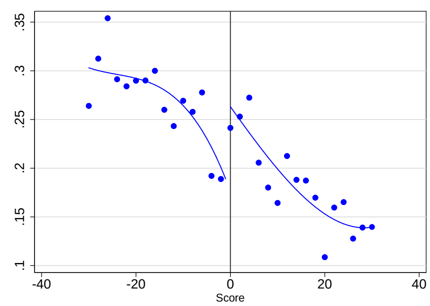

(f) Prob. of CC in Default ‡

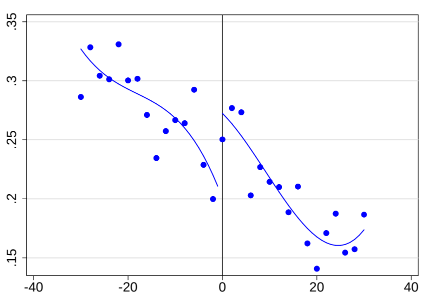

(g) Share of CC in Default †

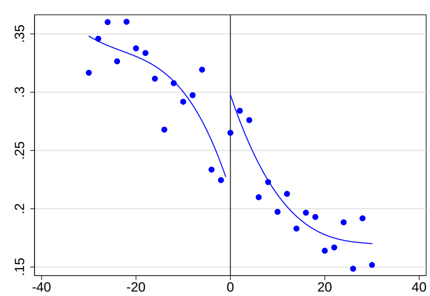

(h) Share of CC in Default ‡

Notes: These figures present the mean of outcome variables regarding long-run (18 months after application) measures of delinquency for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 670 during the application process. Delinquency and default are measured cumulatively from the moment of application up to 18 months after.
† The variable was constructed including all loans that were active at application as well as those opened afterwards.
‡ The variable was constructed including only loans that were active at application.

Figure 10: Effect on Long-Run Delinquency (18 Months) – 700 Sample

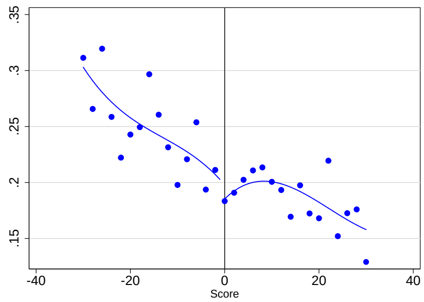

(a) Prob. of CC with 2M Delinq. †

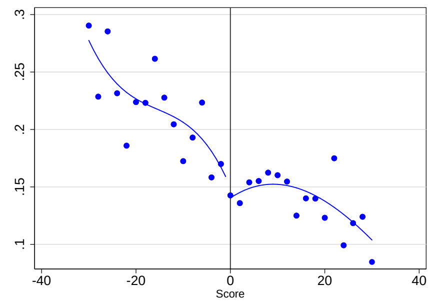

(b) Prob. of CC with 2M Delinq. ‡

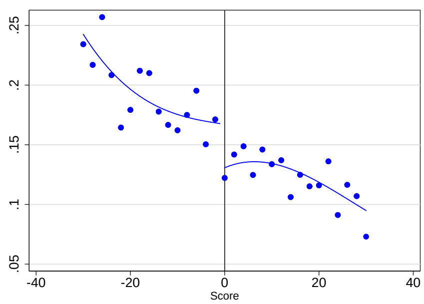

(c) Share of CC with 2M Delinq. †

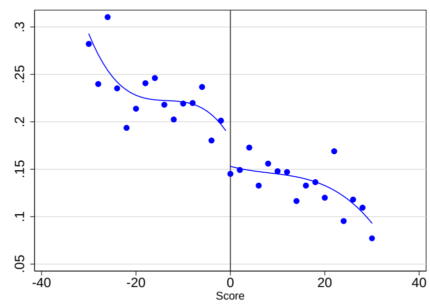

(d) Share of CC with 2M Delinq. ‡

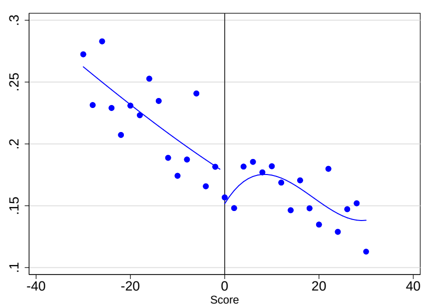

(e) Prob. of CC in Default †

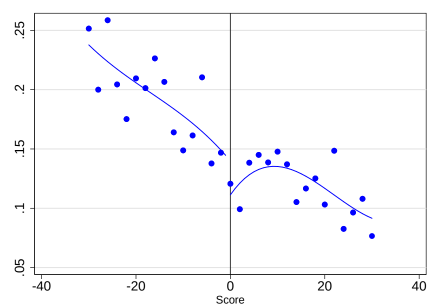

(f) Prob. of CC in Default ‡

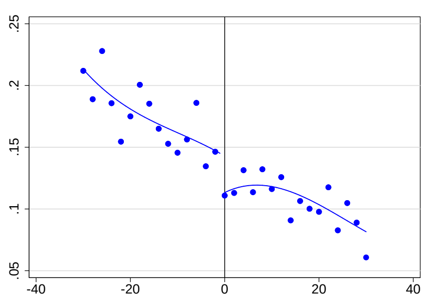

(g) Share of CC in Default †

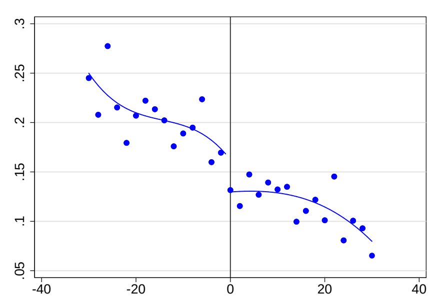

(h) Share of CC in Default ‡

Notes: These figures present the mean of outcome variables regarding long-run (18 months after application) measures of delinquency for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 700 during the application process. Delinquency and default are measured cumulatively from the moment of application up to 18 months after.
† The variable was constructed including all loans that were active at application as well as those opened afterwards.
‡ The variable was constructed including only loans that were active at application.

Table 5: The Effect of Approval on Long-Run Delinquency (18 Months) on Credit Lines Active at the Moment of Application and Other Types of Credit

	Prob. of CC with 2M Delinq. ‡	Share of CC with 2M Delinq. ‡	Prob. of CC in Default ‡	Share of CC in Default ‡	Prob. of Credit Lines in Default Excl. CC †	Share of Credit Lines in Default Excl. CC †	Prob. of Credit Lines in Default Excl. CC ‡	Share of Credit Lines in Default Excl. CC ‡
Panel A: OLS								
Pooled cutoffs	0.014	0.004	0.004	0.003	0.021	0.003	0.030**	0.016
	(0.017)	(0.013)	(0.015)	(0.011)	(0.013)	(0.011)	(0.013)	(0.013)
Above cutoff 670	0.070*	0.071**	0.087**	0.068**	0.099***	0.069**	0.071**	0.067**
	(0.041)	(0.034)	(0.036)	(0.030)	(0.034)	(0.031)	(0.031)	(0.029)
Above cutoff 700	-0.007	-0.018	-0.026*	-0.020*	-0.005	-0.019**	0.016	-0.002
	(0.018)	(0.013)	(0.014)	(0.010)	(0.013)	(0.009)	(0.014)	(0.011)
Panel B: IV								
Pooled cutoffs	0.032	0.010	0.009	0.006	0.047	0.006	0.066**	0.035
	(0.037)	(0.029)	(0.032)	(0.024)	(0.030)	(0.023)	(0.029)	(0.027)
Approved 670	0.150*	0.150**	0.186**	0.145**	0.212***	0.146**	0.151**	0.143**
	(0.085)	(0.070)	(0.073)	(0.059)	(0.081)	(0.070)	(0.065)	(0.060)
Approved 700	-0.015	-0.041	-0.059*	-0.045**	-0.012	-0.043**	0.036	-0.005
	(0.040)	(0.029)	(0.030)	(0.022)	(0.030)	(0.021)	(0.031)	(0.025)
Panel C: Means [-5;-1] from cutoff								
Pooled cutoffs	0.195	0.146	0.170	0.128	0.250	0.117	0.199	0.125
670	0.244	0.186	0.201	0.160	0.353	0.168	0.312	0.188
700	0.179	0.133	0.159	0.118	0.217	0.100	0.162	0.105
N	23492	23492	23492	23492	23492	23492	23492	23492
Panel D: Joint Testing (p-values)								
670 = 700	0.081	0.016	0.002	0.006	0.006	0.011	0.102	0.014

Notes: This table is analogous to Table 4, but focuses on externality effects. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. Measures of delinquency and default are defined in the same way as the variables presented in Table 4, but differ in terms of the types of loans they include. The first four columns include variables constructed including only credit cards that were active at application. The following two columns include all non-credit card loans that were active at application as well as those opened afterwards. The final two columns include only non-credit card loans that were active at application. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure 11: Robustness of ATT Long-Run Results (18 Months)

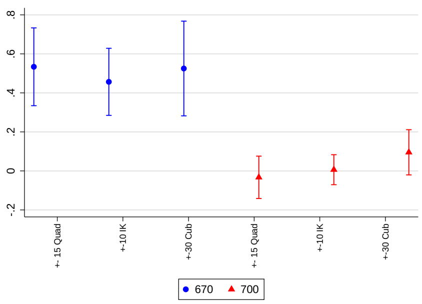

(a) Prob. of CC with 2M Delinq. †

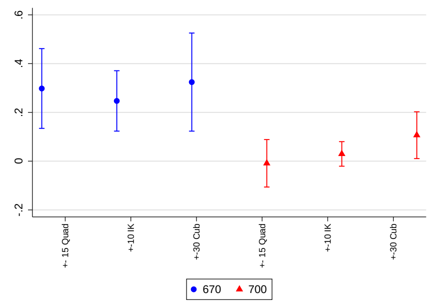

(b) Prob. of CC with 2M Delinq. ‡

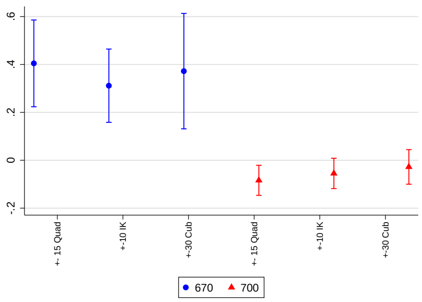

(c) Share of CC with 2M Delinq. †

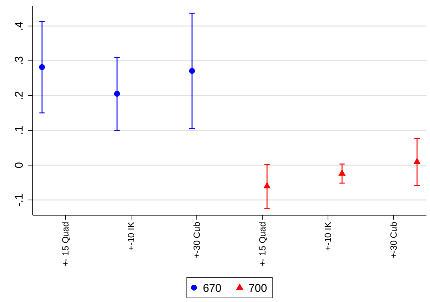

(d) Share of CC with 2M Delinq. ‡

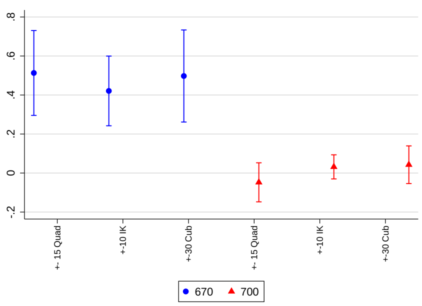

(e) Prob. of CC in Default †

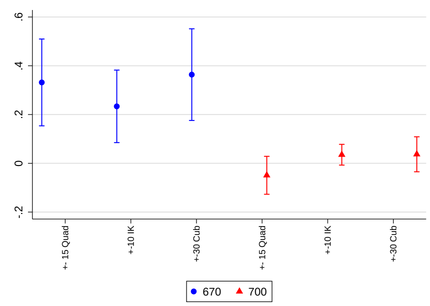

(f) Prob. of CC in Default ‡

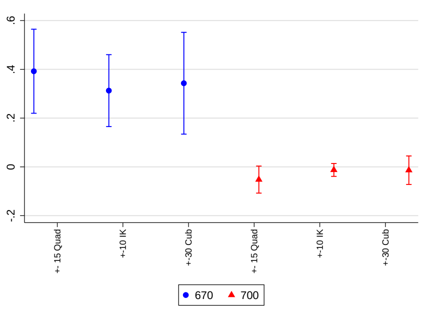

(g) Share of CC in Default †

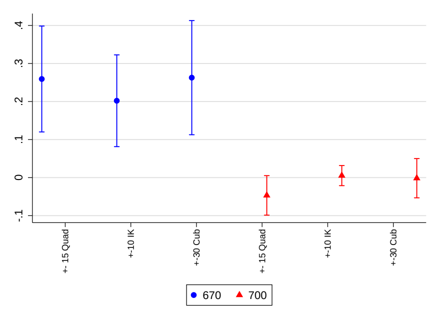

(h) Share of CC in Default ‡

Notes: The figures presents the robustness of the estimated ATT effect of the application being approved on different measures of delinquency and default, using different polynomials (quadratic and cubic), different ranges above the cutoff (10, 15, 60) and those obtained from a local linear regression with optimal bandwidths provided by Imbens and Kalyanaraman (2011). Vertical lines denote 90% confidence intervals (standard errors were clustered at the credit score level). We also included an interaction of the numbers of cards before the bank's decision with the above-the-cutoff indicator variable. Each color represents a different cutoff. Delinquency and default are measured cumulatively from the moment of application up to 18 months after. † The variable was constructed including all loans that were active at application as well as those opened afterwards. ‡ The variable was constructed including only loans that were active at application.

Figure 12: The Effect of Approval on Long-Run Delinquency (18 Months) on Credit Cards Active at the Moment of Application by Number of Months around Change in Cutoff

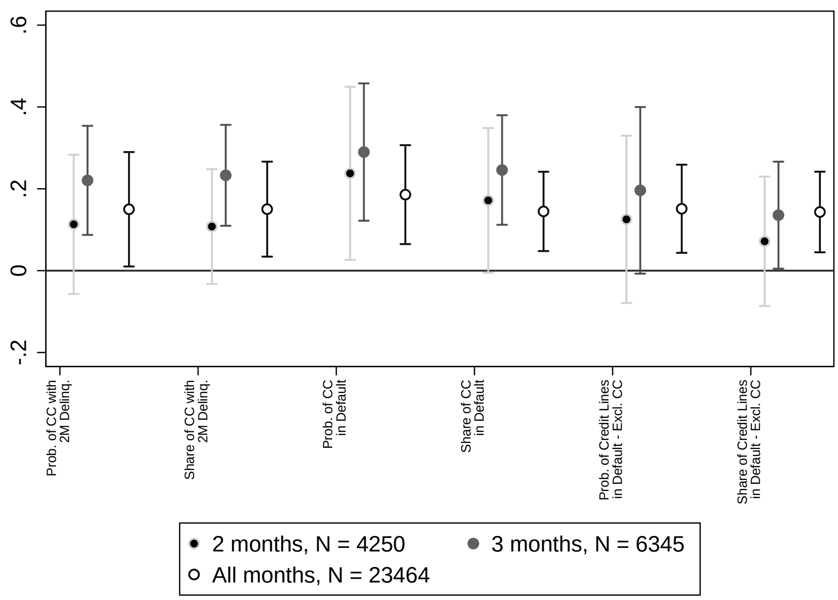

(a) 670 Cutoff

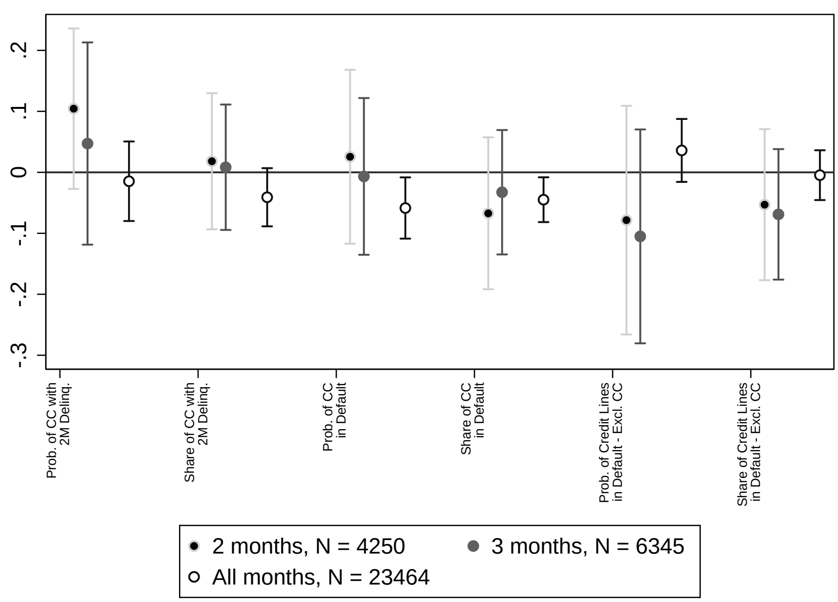

(b) 700 Cutoff

Notes: These figures present the estimated ATT effects for different populations of the 670 and 700 subsamples. The dependent variables were constructed including only loans that were active at application. Panel (a) presents the effects for the 670 group, while Panel (b) presents them for the 700 group. On the horizontal axis both graphs have several measures of default. Delinquency and default are measured cumulatively from the moment of application up to 18 months after. For each cutoff and variable, the figure compares the main ATT results (hollow dots) against estimates obtained using the 3-month (or 2-month) sample, which uses data from applications made in February, March and April 2011 (or March, April 2011, respectively) for the estimates for the 700 group, while it uses data from June, July and August 2011 (or June and July 2011) for the 670 group. Vertical lines denote 90% confidence intervals (standard errors were clustered at the credit score level).

References

Adams, Will, Liran Einav, and Jonathan Levin, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, March 2009, *99* (1), 49–84.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass Through Credit Expansions to Consumers who Want to Borrow?*," *The Quarterly Journal of Economics*, 2017, p. qjx027.

Axelson, Igor and Ulf Makarov, *Sequential Credit Markets*, 2017, (LSE mimeo).

Banxico, "Reporte sobre el sistema financero," Technical Report, Banco de Mexico 2009.

_ , "Reporte sobre las condiciones de competencia en el mercado de emision de tarjetas de credito," *Reporte Banxico*, 2013.

Bisin, Alberto and Danilo Guaitoli, "Moral Hazard and Nonexclusive Contracts," *Rand Journal of Economics*, 2004, *35* (2), 306–328.

Bizer, David and Peter M. DeMarzo, "Sequential Banking," *Journal of Political Economy*, 1992, *100* (1), 41–61.

Bris, Arturo and Ivo Welch, "The optimal concentration of creditors," *The Journal of Finance*, 2005, *60* (5), 2193–2212.

Brunner, Antje and Jan Pieter Krahnen, "Multiple lenders and corporate distress: Evidence on debt restructuring," *The Review of Economic Studies*, 2008, *75* (2), 415–442.

Cappelletti, Giuseppe, Paolo Emilio Mistrulli et al., "Multiple lending, credit lines and financial contagion," Technical Report, Bank of Italy, Economic Research and International Relations Area 2017.

Castellanos, Sara, Diego Jimenez, Aprajit Mahajan, and Enrique Seira, "Mitigating the Risks of Financial Inclusion with contract terms," 2017. Working Paper.

Comision Nacional Bancaria y de Valores, "Reporte de Inclusion Financiera 5," 2013. Comision Nacional Bancaria y de Valores.

Detragiache, Enrica, Paolo Garella, and Luigi Guiso, "Multiple versus single banking relationships: Theory and evidence," *The Journal of Finance*, 2000, *55* (3), 1133–1161.

Dobbie, Will and Marta Skiba, "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending," *American Economic Journal: Applied Economics*, October 2013, *5* (4), 256–282.

Einav, Liran, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf, "Beyond Statistics: The Economic Content of Risk Scores," *American Economic Journal: Applied Economics*, April 2016, *8* (2), 195–224.

Federal Reserve Bank of New York, "Quarterly Report on Household Debt and Credit," 2010. Federal Reserve Bank of New York.

Gobbi, Giorgio and Enrico Sette, "Do firms benefit from concentrating their borrowing? evidence from the great recession," *Review of Finance*, 2013, *18* (2), 527–560.

Gross, David B and Nicholas S Souleles, "An Empirical Analysis of Personal Bankruptcy and Delinquency," *The Review of Financial Studies*, April 2002, *15* (1), 319–347.

Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw, "Evaluating the Effect of an Antidiscrimination Law using a Regression–Discontinuity Design," *NBER Working Paper*, 1999, *7131.*

Imbens, Guido and Karthik Kalyanaraman, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, November 2011, *79* (3).

Imbens, Guido W. and Thomas Lemieux, "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, February 2008, *142* (2), 615–635.

INEGI, "Encuesta nacional de ingresos y gastos de los hogares: ENIGH-2012," 2012.

Karlan, Dean and Jonathan Zinman, "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment," *Econometrica*, November 2009, *77* (6), 1993–2008.

Keys, Benjamin J., Tanmoy K. Mukherjee, Amit Seru, and Vikrant Vig, "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," *Quarterly Journal of Economics*, 2010, *125* (1), 307–362.

Lee, David S. and Thomas Lemieux, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

Lieberman, Andres, "The Value of a Good Credit Reputation: Evidence from Credit Card Renegotiations," *The Journal of Financial Economics*, 2016, *120* (3), 644–660.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Ongena, Steven and David C Smith, "What determines the number of bank relationships? Cross-country evidence," *Journal of Financial intermediation*, 2000, *9* (1), 26–56.

Parlour, Christine A. and Uday Rajan, "Competition in Loan Contracts," *American Economic Review*, 2001, *91* (5), 1311–1328.

Petersen, Mitchell A. and Raghuram G. Rajan, "The Effect of Credit Market Competition on Lending Relationships," *The Quarterly Journal of Economics*, 1995, *110* (2), 407–443.

Thistlethwaite, Donald and Donald Campbell, "Regression–Discontinuity Analysis: An Alternative to the Ex Post Facto Experiment," *Journal of Educational Psychology*, 1960, *51*, 309–317.

US, "Credit Card Market: Economic Benefits and Industry Trends," *International Trade Administration Report, Department of Commerce*, 2008.

APPENDIX

A Merging Procedure and Sample Selection

Our merging procedure starts with data from the entire sample of applications for a Bank A credit card made between January 2010 and September 2012. Out of 604,509 original observations we keep observations with unique identifiers. Furthermore, we only keep the last application made by each individual in case the same individual applied multiple times to Bank A. After this procedure, 484,835 individuals/applications remain in the sample. We then matched this data with the CB data, achieving a 95.5% match (462,842 applicants).

After this match we applied the following criteria. First, since Bank A has a much laxer approval policy with their existing clients (those who have a bank account at the moment of application), no discontinuity in the probability of approval could be exploited. Therefore, to use the RD methodology we were forced to keep only applications from individuals that did not have a bank account in Bank A at the moment of application. Second, during certain months within the sample period Bank A ran several experiments with the credit score threshold that determines eligibility of a new credit card. Thus, in some months there were multiple close credit score thresholds that made the discontinuities in the probability of approval not as strong as those exploited throughout this paper. We drop all applications made within those sub periods. After this selection process, we are left with 106,444 applications, they have credit scores ranging between 400 and 800. Finally, given the local nature of the RD design, we narrowed our final sample to applicants with a credit score (measured at the moment of application) that is within the ± 30 points bounds around the credit score threshold used by Bank A in the approval policy at the relevant threshold regime period.

B Additional Tables and Figures

Figure B1.1: Pre-Treatment Characteristics – 670 Sample

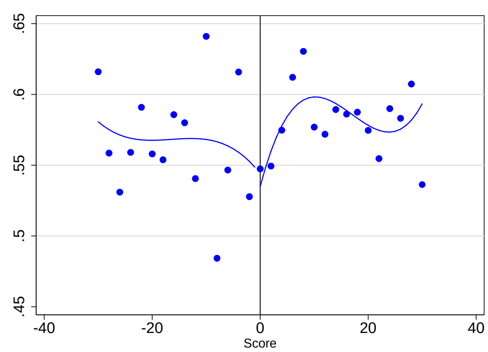

(a) % Male

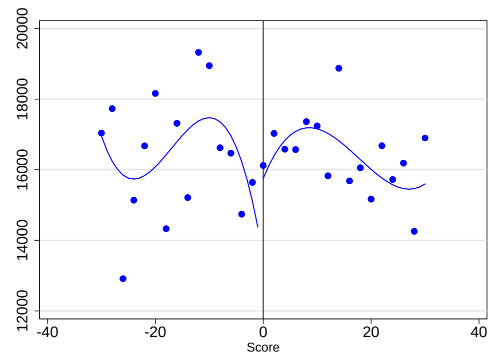

(b) Amount requested (Log)

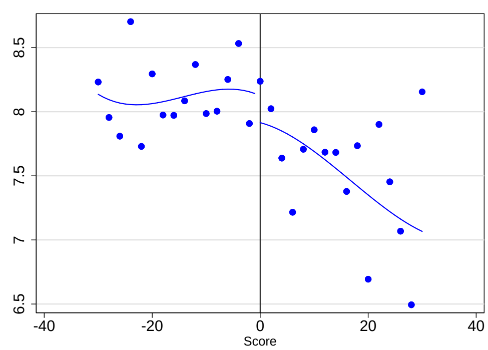

(c) Tenure in Bureau

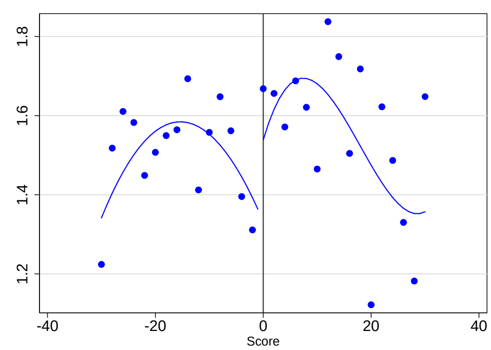

(d) # CC 30 days before

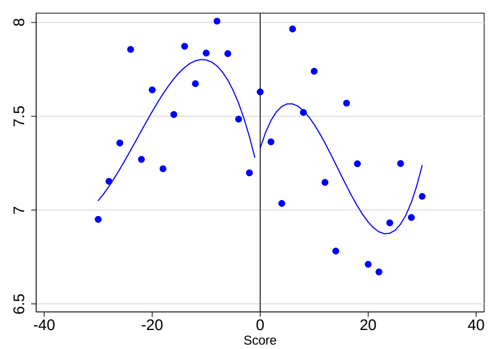

(e) Total debt (Log)

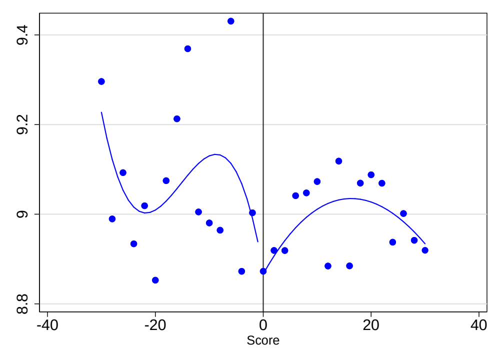

(f) Income (Log)

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 670 during the application process. Panel (a) refers to the percentage of males in each score bin, Panel (b) to the credit limit requested at the application in logs, Panel (c) to the years each person has been in the Credit Bureau, Panel (d) to the number of active credit cards applicants had 30 days before the application, Panel (e) to total Debt in 2010 and panel in logs and Panel (f) to the applicant's administrative income in logs.

Figure B1.2: Pre-Treatment Characteristics – 700 Sample

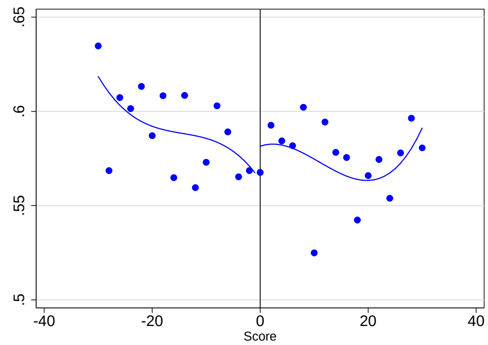

(a) % Male

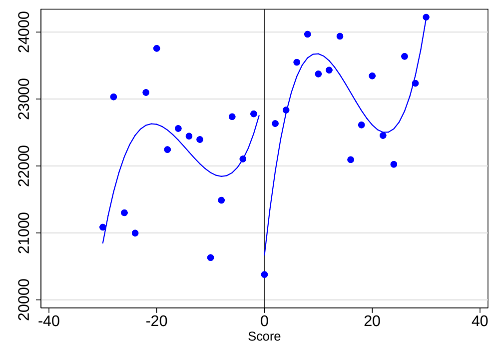

(b) Amount requested (Log)

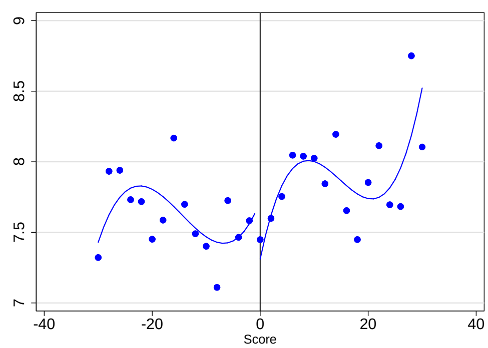

(c) Tenure in Bureau

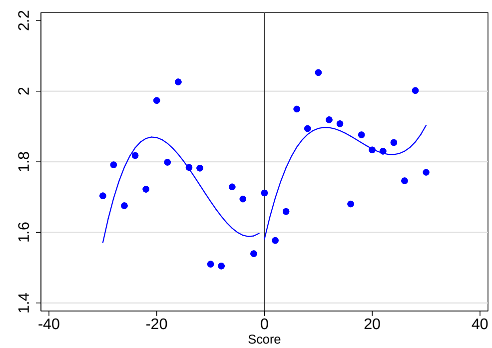

(d) # CC 30 days before

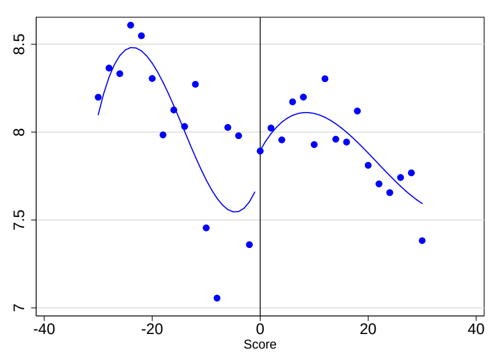

(e) Total debt (Log)

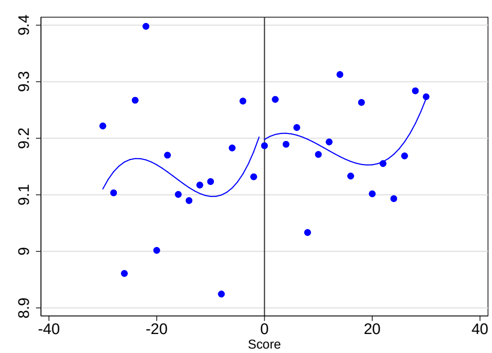

(f) Income (Log)

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 700 during the application process. Panel (a) refers to the percentage of males in each score bin, Panel (b) to the credit limit requested at the application in logs, Panel (c) to the years each person has been in the Credit Bureau, Panel (d) to the number of active credit cards applicants had 30 days before the application, Panel (e) to total Debt in 2010 and panel in logs and Panel (f) to the applicant's administrative income in logs.

Figure B1.3: Pre-Approval Outcome Variables – 670 Sample

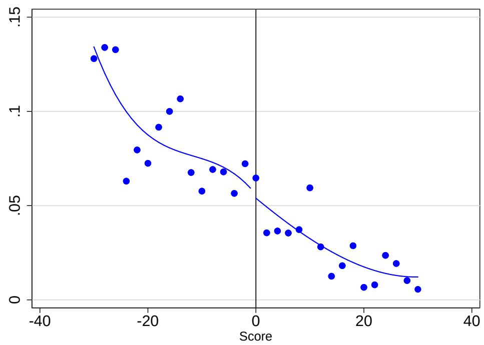

(a) #CC with 2M Delinq.

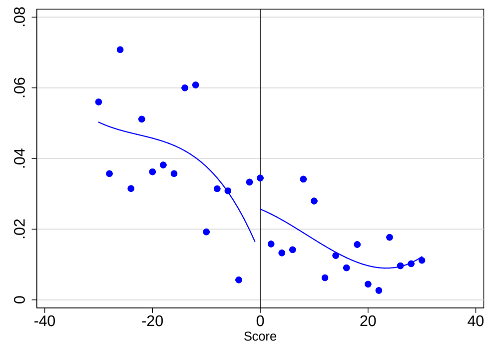

(b) #CC in Default

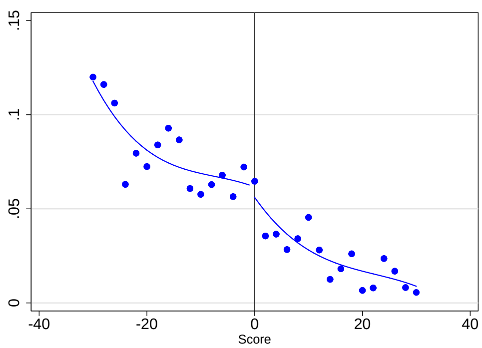

(c) Prob. of CC with 2M Delinq.

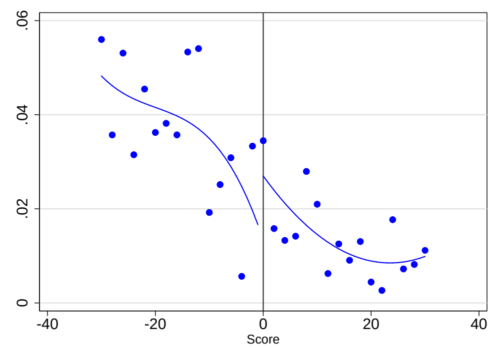

(d) Prob. of CC in Default

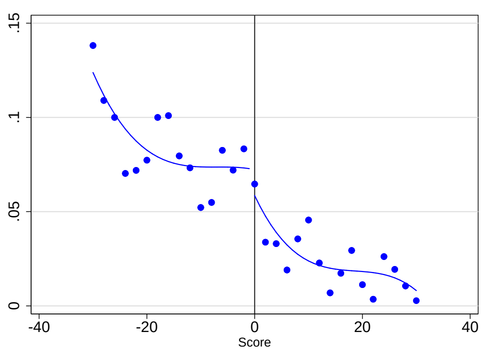

(e) Share of CC with 2M Delinq.

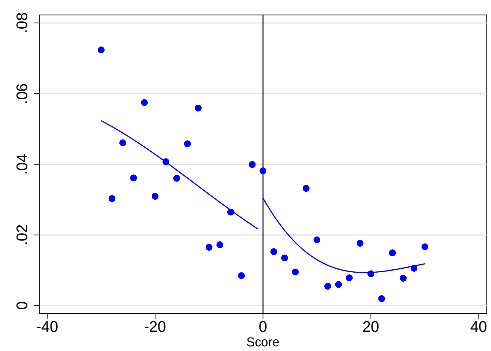

(f) Share of CC in Default

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 670 during the application process. Panel (a) refers to the number of credit cards that had at least one delinquency episode from the earliest month with available information of the card to the application date, Panel (c) to the probability of a delinquency, which is defined as an indicator variable that is equal to one if the applicant has had any credit card with 60 to 90 days past due from the earliest month with available information of the card to the application date, and Panel (e) to ratio of the number of cards in delinquency over the total number of cards. Panels (b), (d) and (f) are analogous but focus on default, which is defined as late payments of 90 days or more. These variables were constructed including only credit cards that were active at the date of application.

Figure B1.4: Pre-Approval Outcome Variables – 700 Sample

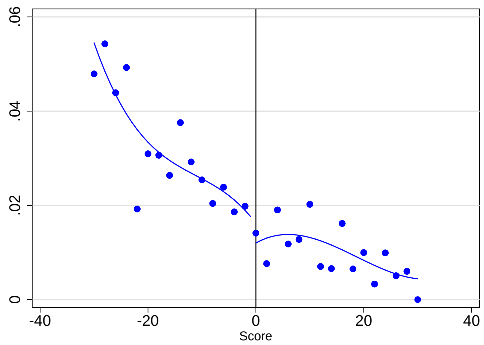

(a) #CC with 2M Delinq.

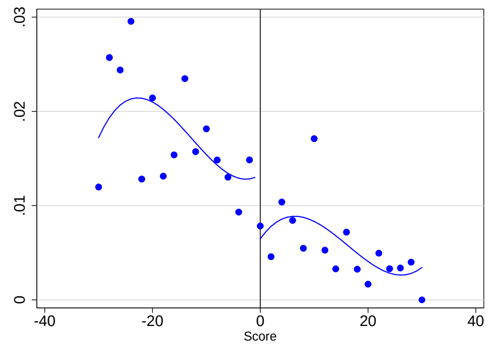

(b) #CC in Default

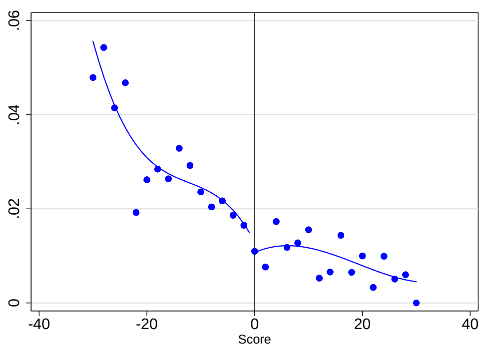

(c) Prob. of CC with 2M Delinq.

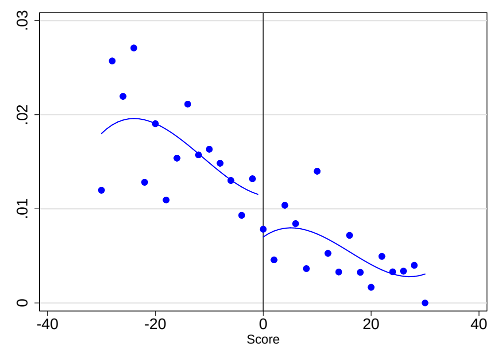

(d) Prob. of CC in Default

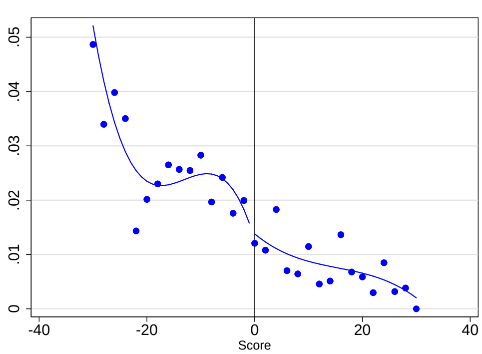

(e) Share of CC with 2M Delinq.

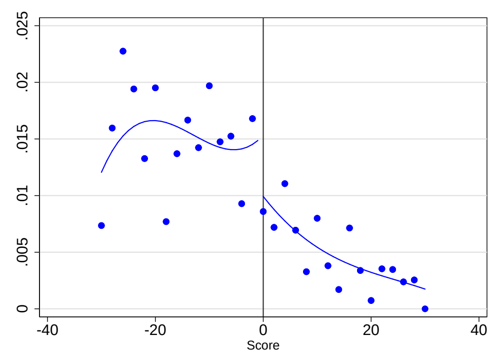

(f) Share of CC in Default

Notes: These figures present the mean of predetermined characteristics for each pair of values of the standardized credit score between standardized scores of -30 and 30. It also displays a polynomial fit of degree 3 to the raw data, allowing the intercept and the coefficients of the polynomial to differ in both sides of the threshold. The vertical line located at 0 represents the cutoff value used by the bank in its assignment process. The sample is restricted to applicants that faced a cutoff of 700 during the application process. Panel (a) refers to the number of credit cards that had at least one delinquency episode from the earliest month with available information of the card to the application date, Panel (c) to the probability of a delinquency, which is defined as an indicator variable that is equal to one if the applicant has had any credit card with 60 to 90 days past due from the earliest month with available information of the card to the application date, and Panel (e) to ratio of the number of cards in delinquency over the total number of cards. Panels (b), (d) and (f) are analogous but focus on default, which is defined as late payments of 90 days or more. These variables were constructed including only credit cards that were active at the date of application.

Table B1.1: The Effect of Approval on Short-Run Delinquency (6 Months) on Credit Lines Active at the Moment of Application

	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	Prob. of CC in Default	Share of CC in Default
	Panel A: OLS			
Pooled cutoffs	-0.020	-0.011	-0.019	-0.009
	(0.015)	(0.013)	(0.013)	(0.012)
Above cutoff 670	0.037*	0.027	0.035**	0.027
	(0.022)	(0.019)	(0.016)	(0.017)
Above cutoff 700	-0.039*	-0.023	-0.038**	-0.021
	(0.022)	(0.018)	(0.015)	(0.014)
	Panel B: IV			
Pooled cutoffs	-0.043	-0.023	-0.042	-0.019
	(0.034)	(0.029)	(0.028)	(0.026)
Approved 670	0.079	0.057	0.075**	0.057*
	(0.051)	(0.042)	(0.034)	(0.034)
Approved 700	-0.088*	-0.052	-0.086**	-0.047
	(0.049)	(0.040)	(0.034)	(0.032)
	Panel C: Means [-5;-1] from cutoff			
Pooled cutoffs	0.099	0.068	0.073	0.050
670	0.127	0.093	0.090	0.065
700	0.090	0.060	0.068	0.046
N	23492	23492	23492	23492
	Panel D: Joint Testing (p-values)			
670 = 700	0.036	0.084	0.000	0.021

Notes: This table reports the RD estimates on different measures of cumulative delinquency and default during the first 6 months after the application. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. The dependent variables were constructed using information only from credit cards that were active at application. Probability of CC with 2M Delinquency is an indicator variable that is equal to one if the applicant has had at least one delinquency episode from the date of application to the first six months after application. Share of CC with 2M Delinquency is the share of cards that were in such a situation during the same period of time. The next two columns are analogous to the former but focus on default, which is defined as late payments of 90 days or more. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table B1.2: Baseline Long-Run (18 Months) Results for Highly Leveraged Sample – All Cards

	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	Prob. of CC in Default	Share of CC in Default
	Panel A: OLS			
Pooled cutoffs	-0.009	-0.023	-0.018	-0.025
	(0.043)	(0.024)	(0.045)	(0.027)
Above cutoff 670	0.230***	0.171**	0.247**	0.152**
	(0.077)	(0.073)	(0.106)	(0.074)
Above cutoff 700	-0.095*	-0.095***	-0.115***	-0.091***
	(0.050)	(0.028)	(0.042)	(0.026)
	Panel B: IV			
Pooled cutoffs	-0.015	-0.038	-0.029	-0.043
	(0.071)	(0.039)	(0.075)	(0.045)
Approved 670	0.356***	0.267**	0.382**	0.237**
	(0.118)	(0.108)	(0.159)	(0.107)
Approved 700	-0.166**	-0.165***	-0.200***	-0.159***
	(0.083)	(0.051)	(0.070)	(0.047)
	Panel C: Means [-5;-1] from cutoff			
Pooled cutoffs	0.304	0.208	0.267	0.190
670	0.321	0.218	0.250	0.176
700	0.298	0.205	0.273	0.195
N	4054	4054	4054	4054
	Panel D: Joint Testing (p-values)			
670 = 700	0.000	0.003	0.001	0.004

Notes: This table reports the RD estimates on different measures of cumulative delinquency and default during the first 18 months after the application. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample is restricted in two ways. First, the sample is restricted to applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. Second, the sample is further restricted to applicants with average debt-to-limit ratio across credit cards in January 2010 above the 70th percentile (i.e., applicants with a credit limit utilization of at least 76%). The dependent variables were constructed using information from all credit cards that were active at application as well as those opened afterwards. Probability of CC with 2M Delinquency is an indicator variable that is equal to one if the applicant has had at least one delinquency episode from the date of application to the first six months after application. Share of CC with 2M Delinquency is the share of cards that were in such a situation during the same period of time. The next two columns are analogous to the former but focus on default, which is defined as late payments of 90 days or more. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table B1.3: Baseline Long-Run (18 Months) Results for Highly Leveraged Sample – Cards and Loans Active at Application

	Prob. of CC with 2M Delinq.	Share of CC with 2M Delinq.	Prob. of CC in Default	Share of CC in Default	Prob. of Credit Lines in Default - Excl. CC	Share of Credit Lines in Default - Excl. CC
Panel A: OLS						
Pooled cutoffs	-0.047	-0.024	-0.044	-0.024	0.018	-0.016
	(0.032)	(0.023)	(0.031)	(0.026)	(0.041)	(0.029)
Above cutoff 670	0.110*	0.175**	0.176**	0.173***	0.102	0.027
	(0.066)	(0.072)	(0.079)	(0.063)	(0.122)	(0.071)
Above cutoff 700	-0.107**	-0.099***	-0.129***	-0.097***	-0.021	-0.034
	(0.044)	(0.032)	(0.035)	(0.029)	(0.039)	(0.027)
Panel B: IV						
Pooled cutoffs	-0.078	-0.040	-0.074	-0.040	0.031	-0.026
	(0.052)	(0.037)	(0.050)	(0.042)	(0.068)	(0.047)
Approved 670	0.170*	0.272**	0.273**	0.268***	0.158	0.042
	(0.102)	(0.112)	(0.122)	(0.089)	(0.177)	(0.106)
Approved 700	-0.186***	-0.172***	-0.224***	-0.169***	-0.037	-0.060
	(0.072)	(0.056)	(0.059)	(0.052)	(0.067)	(0.044)
Panel C: Means [-5;-1] from cutoff						
Pooled cutoffs	0.276	0.205	0.245	0.188	0.178	0.106
670	0.298	0.202	0.238	0.173	0.238	0.147
700	0.269	0.206	0.248	0.193	0.157	0.091
N	4054	4054	4054	4054	4054	4054
Panel D: Joint Testing (p-values)						
670 = 700	0.017	0.004	0.001	0.001	0.353	0.387

Notes: This table is analogous to Table 4, but focuses on externality effects for applicants highly leveraged. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample is restricted in two ways. First, the sample is restricted to applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. Second, the sample is further restricted to applicants with average debt-to-limit ratio across credit cards in January 2010 above the 70th percentile (i.e., applicants with a credit limit utilization of at least 76%). The first four columns include dependent variables constructed including only credit cards that were active at application. The final two columns include only non-credit card loans that were active at application. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure B1.5: The Effect of Approval on Long-Run Delinquency (18 Months)
by Number of Active Credit Cards at Application

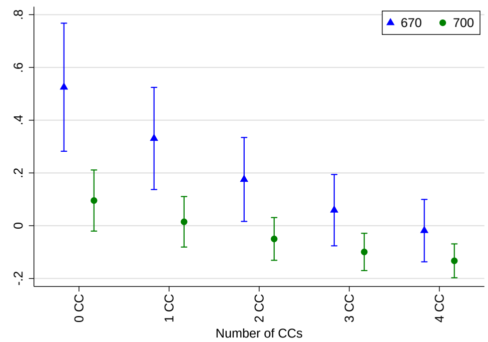

(a) Prob. of CC with 2M Delinq.

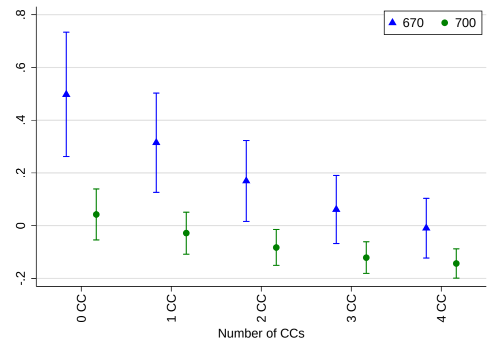

(b) Prob. of CC in Default

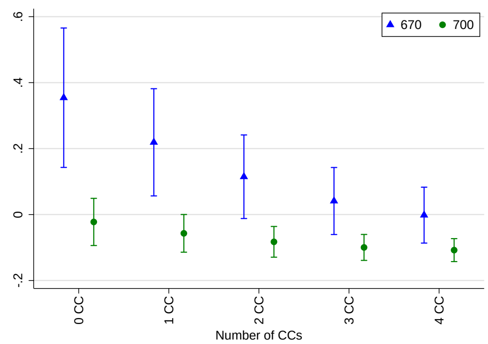

(c) Share of CC with 2M Delinq.

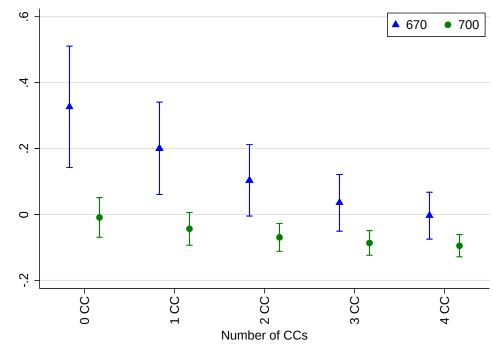

(d) Share of CC in Default

Notes: These figures show the regression discontinuity estimates of the effect of the application being approved on different measures of delinquency by the number of credit cards active at application. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. Probability of CC with 2M Delinquency is an indicator variable that is equal to one if the applicant has had at least one delinquency episode from the date of application to the first six months after application. Share of CC with 2M Delinquency is the share of cards that were in such a situation during the same period of time. The next two columns are analogous to the former but focus on default, which is defined as late payments of 90 days or more. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. In order to allow for heterogeneous effects, we included an interaction of the numbers of active cards before the bank's decision with the above-the-cutoff indicator variable. These figures present the predicted effects (and their 95% confidence intervals) for applicants with a different number of credit cards at the moment of application. Each color represents the estimation for a different sample. Clustered standard errors at the credit score level are reported in parenthesis.

Table B1.4: The Effect of Additional Credit Limit on Long-Run Delinquency (18 Months)

	Prob. of CC in Default All cards	Share of CC in Default All Cards	Prob. of CC in Default Cards open at app.	Share of CC in Default Cards open at app.
		Panel A: IV		
Approved Amount 670	0.015***	0.010**	0.011**	0.008**
	(0.006)	(0.004)	(0.004)	(0.004)
Approved Amount 700	-0.004	-0.004**	-0.004*	-0.003**
	(0.003)	(0.002)	(0.002)	(0.002)
		Panel B: Means [-5;-1] from cutoff		
670	0.24	0.17	0.20	0.16
700	0.19	0.13	0.16	0.12
N	23492	23492	23492	23492
		Panel C: Joint Testing (p-values)		
670 = 700	0.000	0.001	0.001	0.003

Notes: This table presents OLS estimates of the RD specification $y_{it} = \alpha + \beta\gamma ApprovedAmount_i + f(score_{it}, \nu^-, \nu^+) + X'\xi + \nu_{it}$, where $ApprovedAmount_i$ is instrumented with the threshold dummy $\mathbf{1}(score_{it} \geq \overline{score}_t)$. Panel A presents the IV results for each subsample. Panels B displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel C presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. The first two columns include all credit cards that were active at application as well as those opened afterwards. The last two columns include only credit cards that were active at application. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table B1.5: Effects on Debt

	Current Debt New Bank A CC	Prob. CC Debt > 50th perc.	Prob. Non-CC Debt > 50th perc.	Prob. CC Debt > 75th perc.	Prob. Non-CC Debt > 75th perc.
Panel A: OLS					
Pooled cutoffs	998.986***	0.067***	0.007	0.014	0.016
	(116.440)	(0.022)	(0.020)	(0.019)	(0.015)
Above cutoff 670	1525.943***	0.105***	-0.002	0.054**	0.083***
	(338.072)	(0.039)	(0.052)	(0.027)	(0.021)
Above cutoff 700	798.552***	0.053**	0.012	-0.003	-0.005
	(151.084)	(0.025)	(0.025)	(0.025)	(0.018)
Panel B: IV					
Pooled cutoffs	2250.948***	0.150***	0.017	0.031	0.037
	(227.627)	(0.046)	(0.045)	(0.043)	(0.034)
Approved 670	3295.717***	0.228***	-0.004	0.117*	0.179***
	(421.009)	(0.078)	(0.111)	(0.064)	(0.045)
Approved 700	1829.201***	0.120**	0.028	-0.007	-0.012
	(319.680)	(0.057)	(0.057)	(0.057)	(0.041)
Panel C: Means [-5;-1] from cutoff					
Pooled cutoffs	38.836	0.446	0.505	0.225	0.254
670	0.215	0.385	0.523	0.160	0.235
700	51.299	0.465	0.499	0.245	0.260
N	24180	24180	24180	24180	24180
Panel D: Joint Testing (p-values)					
$670 = 700$	0.084	0.215	0.827	0.177	0.001

Notes: This table reports the RD estimates of the effect of eligibility on different measures of debt. Panel A presents the OLS results for each subsample, while Panel B presents the IV results for each subsample. Panels C displays the mean of the dependent variable for applicants with standardized credit scores 5 points below the cutoff. Finally, Panel D presents the p-value of the test of the null hypothesis that the magnitude of the discontinuity is the same across samples. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. In column (1) the dependent variable is the total debt (in MXN) at the moment of the snapshot in the new credit card approved by Bank A. The average debt is very low in the $[-5,-1]$ subsample, since by construction most applicants in that subsample did not receive the new credit card from Bank A. In column (2) the dependent variable is the probability that an applicant's total credit card debt is above the unconditional median. Column (3) is similarly defined, but focuses on debt in non-credit card loans. The dependent variables in the last two columns are similarly defined, but the thresholds are increased to the 75th percentile. All regressions control for a third order polynomial, allowing for a discontinuity of the standardized score at the value of 0. Regressions also include as control variables a set of indicator variables for each month during the application period and for each number of credit cards and other types of loans active at the moment of the application. Clustered standard errors at the credit score level are reported in parenthesis. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure B1.6: Robustness of ITT Long-Run Results (18 Months)

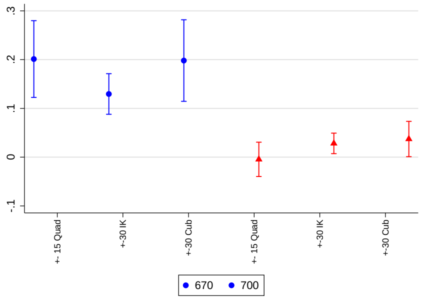

(a) Prob. of CC with 2M Delinq. †

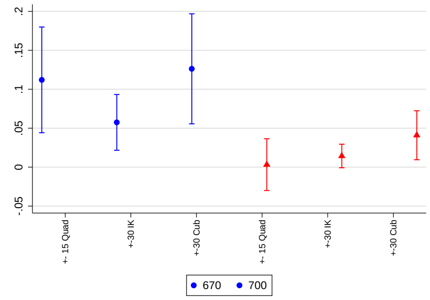

(b) Prob. of CC with 2M Delinq. ‡

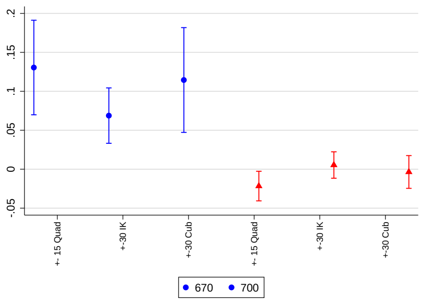

(c) Share of CC with 2M Delinq. †

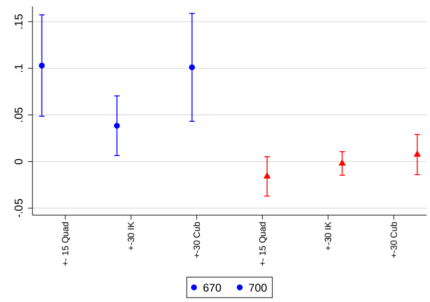

(d) Share of CC with 2M Delinq. ‡

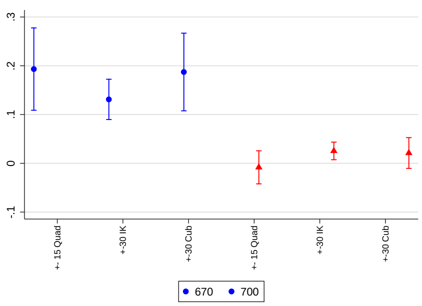

(e) Prob. of CC in Default †

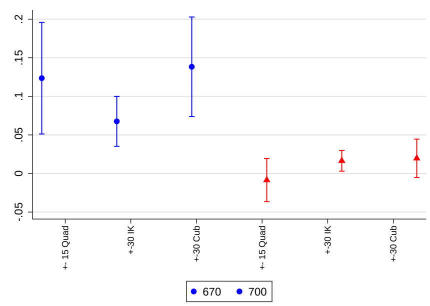

(f) Prob. of CC in Default ‡

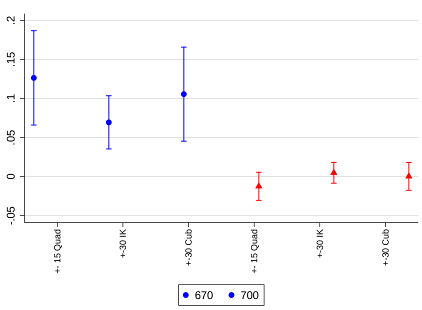

(g) Share of CC in Default †

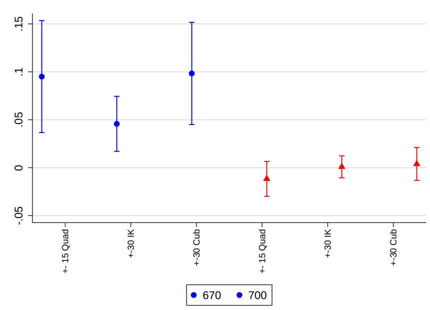

(h) Share of CC in Default ‡

Notes: The figures presents the robustness of the estimated ITT effect on different measures of delinquency and default, using different polynomials (quadratic and cubic), different ranges above the cutoff (10, 15, 60) and those obtained from a local linear regression with optimal bandwidths provided by Imbens and Kalyanaraman (2011). Vertical lines denote 90% confidence intervals (standard errors were clustered at the credit score level). We also included an interaction of the numbers of cards before the bank's decision with the above-the-cutoff indicator variable. Each color represents a different cutoff. Delinquency and default are measured cumulatively from the moment of application up to 18 months after. † The variable was constructed including all loans that were active at application as well as those opened afterwards. ‡ The variable was constructed including only loans that were active at application.

Figure B1.7: The Effect of Approval on Long-Run Delinquency (18 Months) by Number of Months around Change in Cutoff

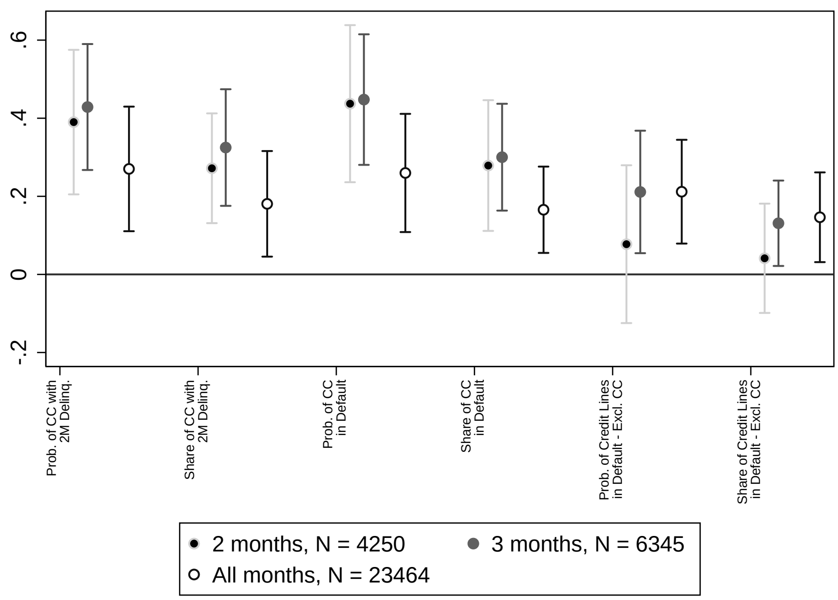

(a) 670 Cutoff

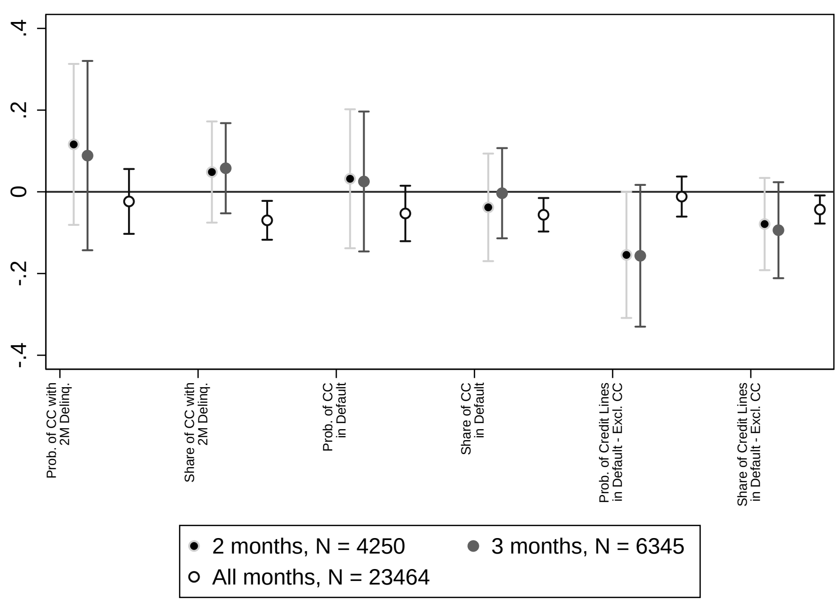

(b) 700 Cutoff

Notes: These figures present the estimated ATT effects for different populations of the 670 and 700 subsamples. The dependent variables were constructed including loans that were active at application as well as those opened afterwards. Panel (a) presents the effects for the 670 group, while Panel (b) presents them for the 700 group. On the horizontal axis both graphs have several measures of default. Delinquency and default are measured cumulatively from the moment of application up to 18 months after. For each cutoff and variable, the figure compares the main ATT results (hollow dots) against estimates obtained using the 3-month (or 2-month) sample, which uses data from applications made in February, March and April 2011 (or March, April 2011, respectively) for the estimates for the 700 group, while it uses data from June, July and August 2011 (or June and July 2011) for the 670 group. Vertical lines denote 90% confidence intervals (standard errors were clustered at the credit score level).

Figure B1.8: Evolution of Average Applicants' Characteristics

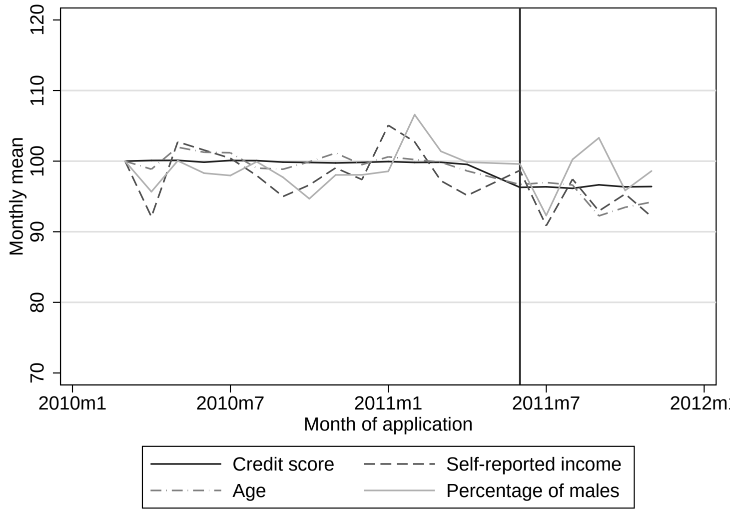

Notes: This figure presents the evolution of the average of applicants' characteristics. Each series has been normalized to its corresponding values as of the first month available in our applications data. The vertical line marks the month in which Bank A started using 670 as the threshold value in the approval process.

Figure B1.9: Comparison of Default Rates between Applicants with Scores below the Threshold across Experiments

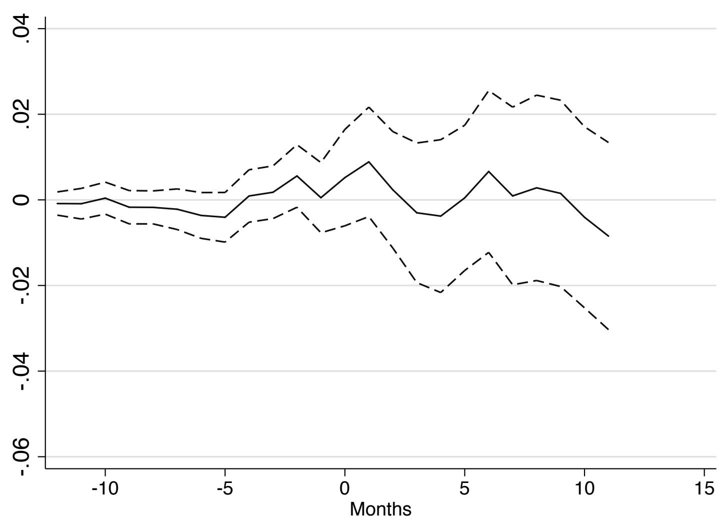

Notes: This figure plots the coefficient of the regression $y_{it} = \alpha_t + \beta_t Cutoff_i^{670} + \epsilon_{it}$, where the dependent variable is the indicator that is equal to one if applicant i is in default in period t (normalized as months-after-application), and $Cutoff_i^{670}$ indicates whether applicant i applied during the 670-threshold regime. β_t captures the probability of being in default t months before/after application for applicants that applied during the 670-regime relative to those that applied during the 700-regime. The sample consists of applicants that had a score between 640 and 660 (this restriction yields 4,315 observations). The figure reports the estimates obtained by running the regression for each t, together with the 90% confidence interval. Standard errors are clustered at the credit score level.

Table B1.6: Default allocation

	Default 3m		Default 6m		Default 12m		Mean
Biggest loan	0.00404		0.00604		0.00806		0.173
	(0.00959)		(0.00964)		(0.00967)		
Oldest loan	0.0308***		0.0191**		0.0108		0.174
	(0.00923)		(0.00928)		(0.00932)		
5 biggest banks	0.0593**	0.0509**	0.0386	0.0293	0.0416	0.0282	0.0245
	(0.0259)	(0.0258)	(0.0261)	(0.0260)	(0.0262)	(0.0261)	
Secured loan	-0.0632***	-0.0415***	-0.0591***	-0.0373**	-0.0740***	-0.0508***	0.0787
	(0.0153)	(0.0152)	(0.0153)	(0.0153)	(0.0154)	(0.0153)	
Ln(credit limit)		-0.00119		-0.00116		-0.00202	8.032
		(0.00137)		(0.00138)		(0.00138)	
Ln(months old)		0.0327***		0.0273***		0.0199***	2.919
		(0.00385)		(0.00387)		(0.00389)	
Constant	0.298***	0.217***	0.351***	0.285***	0.394***	0.355***	
	(0.00476)	(0.0167)	(0.00478)	(0.0168)	(0.00480)	(0.0169)	
N	17041	16920	17041	16920	17041	16920	
Mean Dep. Var.	0.297	0.295	0.346	0.343	0.386	0.382	
R^2	0.235	0.246	0.286	0.297	0.314	0.323	

Notes: This table reports regression results that illustrate the applicants' decision to which loan default. The sample consists of all applicants with standardized credit score at most 30 points above or 30 points below of their respective cutoff value. The sample is further restricted to applicants who had their credit card application approved and who also defaulted on at least one loan after the application. Regressions are conducted using observations at the loan-level, and include observations from loans that were active when the first default episode occurred. The regression estimates the correlation between a card being in default within the first 3, 6 and 12 months after the first default occurred (observations include the first default) and different loan characteristics: loan size, loan seniority, belonging to the 5 biggest banks and being a secured loan. These regressions also include applicant fixed effects.