Series Name: WPS
Paper No.: 151
Issue Date: 5 Mar 2021

Too Fast, Too Furious? Digital Credit Delivery Speed and Repayment Rates

Alfredo Burlando, Michael A. Kuhn and Silvia Prina

Too Fast, Too Furious?

Digital Credit Delivery Speed and Repayment Rates

Alfredo Burlando[§] Michael A. Kuhn[¶] and Silvia Prina[‖]

Abstract

Digital loans are a source of fast short-term credit for millions of people. While digital credit broadens market access and reduces frictions, default rates are high. We study the role of the speed of delivery of digital loans on repayment. Our study uses unique administrative data from a digital lender in Mexico and a regression-discontinuity design. We show that reducing loan speed by doubling the delivery time from ten to twenty hours decreases the likelihood of default by 21%. Our finding hints at waiting periods as a potential consumer protection measure for digital credit.

JEL Classifications: D14, D18, G51, O16

Keywords: Digital credit, waiting periods, defaults, financial access

[§]University of Oregon and CEGA. Department of Economics, 1285 University of Oregon, Eugene, OR 97403. E-mail: burlando@uoregon.edu

[¶]University of Oregon and CEGA. Department of Economics, 1285 University of Oregon, Eugene, OR 97403. E-mail: mkuhn@uoregon.edu

[‖]Northeastern University, Department of Economics, 310A Lake Hall, 360 Huntington Avenue, Boston, MA 02115, United States. E-mail: s.prina@northeastern.edu.

1 Introduction

The digital credit market has recently emerged as a source of fast, automated, remotely provided, short-term loans for millions of people in low- and middle-income countries (Francis, Blumenstock, and Robinson, 2017). Data harvesting and analytics have enabled digital credit providers to assess consumers' creditworthiness and ability to repay without requiring any collateral to secure loans (Björkegren and Grissen, 2018). Thus, digital credit has the potential to help households cope with unexpected shocks and reduce liquidity constraints for investments (e.g., Karlan and Zinman, 2010; Morse, 2011). Indeed, Bharadwaj, Jack, and Suri (2019) find that digital credit in Kenya has improved household resilience to negative shocks. Furthermore, the fast speed of loan provision allows borrowers to act on time-sensitive opportunities to a much greater degree than in the past.

While the speed and ease of access to digital credit makes these loans very appealing, many borrowers struggle to repay them (Carlson, 2017). Digital credit can exacerbate self-control problems, causing overindebtedness and default (Skiba and Tobacman, 2019), making it harder to pay bills (Melzer, 2011) and reducing access to future loans if defaulters are reported to a credit bureau (as is the case in our study).[1] In addition, anecdotal evidence shows that borrowers do not fully understand the terms of their loans (e.g., Mazer and Fiorillo, 2015; McKee, Kaffenberger, and Zimmerman, 2015) and may use them to finance unproductive, time-sensitive investment and consumption opportunities such as gambling (Malingha, 2019). This is particularly important, given that the industry suffers from high default rates (27% in our data). Hence, it is not surprising that policy makers have started to advocate for consumer protection measures targeting the digital credit market (Donovan and Park, 2019).

In this paper, we study the role of the speed of delivery of digital loans on repayment.

[1]Evidence from the credit card market shows that less-sophisticated borrowers may be susceptible to overborrowing, penalties, and backloading repayments (Meier and Sprenger, 2010; Heidhues and Kőszegi, 2010).

To date, this policy-relevant issue has not been studied, despite the continuous growth of this market. We address this knowledge gap with a unique administrative dataset of digital loans and quasi-experimental variation in the time it takes for a loan to be deposited into the borrower's bank account. Specifically, our data consist of loan records from the full set of approved clients from a digital lender operating in Mexico over a seven-month period in 2018-2019. These records include both loan application timestamps and disbursement timestamps, which we use to measure loan delivery speeds. The quasi-experimental variation in delivery speeds comes from the fact that the company disburses loans in batches, a process that occurs only two to four times during the day. Loans added first to a new batch remain in the batch longer than those added last, leading to systematic differences in processing times between loans. Our empirical strategy identifies the discontinuous changes in processing times created each time an existing batch is disbursed and a new batch is opened. Crucially, disbursement times are ex ante unknown to borrowers, and they change daily. Thus, there is no concern that clients can time their applications for faster service. However, unlike the standard regression discontinuity (RD) setup, we do not observe the precise moment a batch is closed; we construct proxies for these cutoff times using a machine-learning technique applied to our disbursement and application submission time data.

On average, for all borrowers, loans submitted just after one of these proxied cutoffs face an additional delay of 9.81 hours, roughly doubling the total amount of time it takes to obtain a loan. We find that the delay induced by missing a batch cutoff increases repayment by 5.6 percentage points, corresponding to a roughly 8% increase relative to similar loans that do not experience the extra delay. These point estimates translate to a 21% reduction in the likelihood of loan default when loan delivery is slowed down. This finding is in line with estimates found in other types of financial market interventions within the microfinance literature.[2] While our ability to explore mechanisms is constrained by the sparse administrative

[2]The study closest to ours, Karlan and Zinman (2009), finds a 2.5 p.p. reduction in loans in collection status when borrowers are offered dynamic incentives (a 21% reduction). Field, Pande, Papp, and Rigol

data we have access to, suggestive evidence points to behavioral biases and intrahousehold bargaining as potential mechanisms of these effects.

Our results are related to recent studies in economics showing that waiting periods without any choice restrictions can affect behavior (Imas, Kuhn, and Mironova, 2016; DeJarnette, 2018; Brownback, Imas, and Kuhn, 2019; Thakral and Tô, 2020). Waiting periods are already used in settings in which myopia and impulsivity are perceived to be particularly harmful. For example, many U.S. states require waiting periods prior to the purchase of firearms (Koenig and Schindler, 2018; Edwards, Nesson, Robinson, and Vars, 2018). They are also implemented in negotiations (Brooks, 2015) and conflict resolution (Burgess, 2004). Our study also relates to the more traditional literature on behavioral biases in consumer financial choice. Behavioral biases induce agents to engage in suboptimal behavior, such as reducing earnings from investments (e.g., Duflo, Kremer, and Robinson, 2011; Kremer, Lee, Robinson, and Rostapshova, 2013) or reducing savings (Dupas and Robinson, 2013). A common solution to these biases is to design financial products that impose restrictions on agents.[3]

2 Setting

Our sample consists of loans from an online digital lender in Mexico. The loan amounts range from 1,500 to 3,000 Mexican pesos (approximately USD 75 to 150),[4] and the loan terms vary from seven to 30 days. The APRs reach up to 478.8%. The characteristics of

(2013) finds that providing a repayment grace period reduces repayments by 6 p.p. (a 370% change relative to mean default); Feigenberg, Field, and Pande (2013) varies MFI group meeting intensity and finds that more frequent meetings increase repayments by 5.1 p.p. (a 72% decrease in mean default); Karlan, Morten, and Zinman (2015) likewise finds a 3.7 p.p. decrease in loans with unpaid balance after 30 days when borrowers are sent SMS reminders (a 27% reduction).

[3]Examples include commitment savings accounts (Ashraf, Karlan, and Yin, 2006), and frequent fixed payments for microfinance borrowers (Bauer, Chytilová, and Morduch, 2012; Field et al., 2013).

[4]The exchange rate during the study period is approximately USD 1= MXP 20.

this loan product are comparable to those of other digital lenders in the market. Potential borrowers interact with the lender using a browser on a smartphone or a computer. The lender's home page prominently reports the interest rate and other costs, including taxes and fees, at the bottom of the window. Potential borrowers are advised that they can get a loan in "minutes."

2.1 Loan application and delivery process

Users start their application by selecting the amount and term of the loan. Applicants need to satisfy the following requirements to obtain a loan: proof of citizenship (a photo of the national identification card); age between 20-65 years; a photo taken from a phone or computer camera; regular income (from a credit report); cellphone number and e-mail address; and a bank account. For first-time applicants the digital lender pulls the applicant's credit history from a credit bureau.

Loan application and preapproval occur online during a single browsing session. Successful applicants are notified that their loans have been preapproved and will be issued once they have been processed. Borrowers undergo verification, which, for first-time clients, includes a call from a customer service representative.

Processed loans are entered into a spreadsheet, which serves as a delivery queue. Loans accumulate in the queue until an employee sends the whole batch to the lender's bank for processing. Once the bank receives a batch, all loans in the batch are disbursed immediately to borrowers' bank accounts. Loans can be repaid anytime after they have been deposited, but the repayment amount includes interest for the full approved duration of the loan.

2.2 Sample

Our sample consists of 11,512 approved loan applications from 7,206 unique borrowers, with loans disbursed between November 2018 and May 2019.[5] Forty-eight percent of the loans in our sample are from first-time borrowers. For any borrower, we observe up to three loans. We are given access to the following administrative data: the timestamps of all loan application submissions and loan disbursements; the repayment status and date of final repayment for each loan; the borrower's age, sex, marital status, number of dependents, and personal income as reported in their first loan application; and the loan sequence (whether this is the borrower's first, second, or third loan). Furthermore, we have information on requested and approved loan amounts and terms for first-time loans but not for repeat loans.

As shown in Appendix Table A1, the borrowers are poorer than the average Mexican worker, with a self-reported median monthly income below 1,000 pesos (52 USD). Of all the borrowers, 45% are female, and 11% lack a credit report. On average, first-time borrowers receive 1,785 pesos (approximately 25% of the average monthly income). Loan processing times, which we refer to as delays, are calculated as the time difference between loan application submission by the client and disbursal by the bank. On average, first-time borrowers face a delay of 26 hours, while for repeat borrowers, it is 9 hours.

Our main outcome variable is repayment. On average, 73.3% of the loans in our sample are repaid, which implies a default rate of 26.7%. For first-time loans. the default rate is 32%, while for repeat borrowers, it is 22%.[6] A lower rate for the latter group is expected since repeat loans are given conditional on past repayment. Appendix Table A2 shows the relationship between borrower/loan characteristics and repayment likelihood. As expected,

[5]The raw data from the lender contain 15,882 loans. Of these loans, 669 had missing submission times, and three were reported disbursed before they were submitted. Sections 3.1 and 3.2 detail the additional steps to determine the estimation sample.

[6]Unfortunately, we cannot tell whether overdue loans have been partially repaid. It is possible that some of the defaulted loans were repaid after we received the data.

income and credit score tend to positively correlate with repayment. The term of a loan correlates negatively with repayment, but the amount of the loan does not.

3 Empirical strategy

Our empirical strategy takes advantage of the fact that while loan applications occur continuously throughout the day, loans are disbursed in batches. We compare loans that are submitted by clients in time to be included in a particular batch to those submitted slightly later that do not. Crucially, borrowers are unaware of this batching process. In addition, in any given day, there are no set times at which batches are sent to the bank for disbursal.[7]

Figure 1 shows a simplified timeline of loan applications and disbursements to illustrate our approach to identification. Individuals apply for loans at different points in time. Once processed, loans are assigned to the existing disbursement batch. For example, loans k and l are both processed prior to the Batch A cutoff, and thus assigned to Batch A and disbursed at t_2, while loans m, n, and o are approved after Batch A has been disbursed. Thus, they are assigned to Batch B and disbursed at t_4.

For each batch, we define a batch cutoff as the latest point in time at which a loan application could be submitted by a client and make it into that batch. This means that no loans received after a batch cutoff can possibly be in that batch. However, it is also possible that some loans received prior to a cutoff will end up in later batches. For example, in Figure 1, both loans l and o are submitted prior to the Batch A cutoff. Loan l is quickly approved and ends up in Batch A, while loan o takes longer and ends up in Batch B.

Our empirical strategy is best illustrated by the comparison between loans l and m. These loans have been submitted by two separate clients around the same time and take a similar amount of time to be verified. However, because they fall on different sides of the Batch A cutoff time t_1, loan l is delivered much more quickly.

[7]This situation also implies that the lender is not aware of these batch cutoff times in advance, either.

Figure 1: **Hypothetical timeline of loan submission, verification and disbursement**

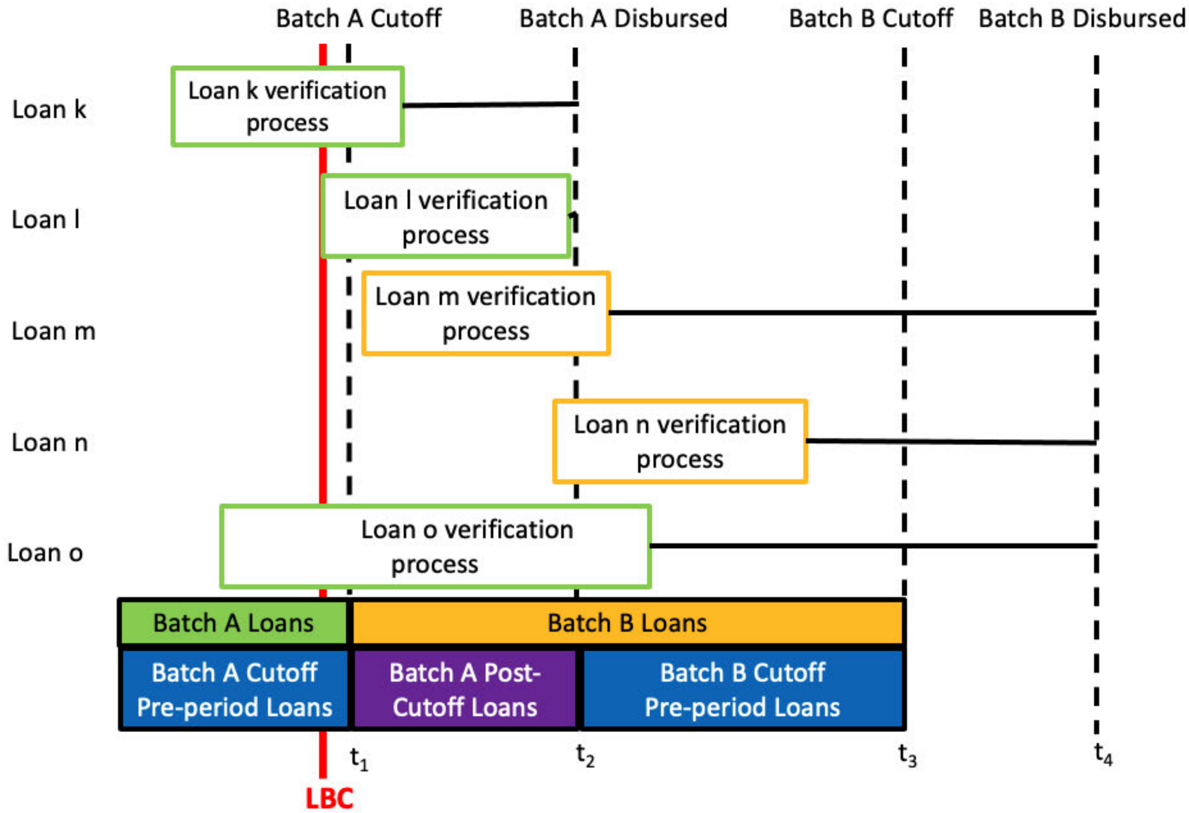

Loan verification process includes the time between application submission and pre-approval by the client and the time placement of the approved loan into the loan delivery queue (the batch). The LBC line stands for "lower bound cutoff", as defined in section 3.1.

To implement this strategy, we first assign every loan to the closest batch cutoff (based on its application submission time). Next, we create an indicator called $PostBatch$ that takes a value of one if the application was submitted after its assigned cutoff. In our example, the indicator takes the value of zero for loan l and one for loan m. Then, we compute a continuous variable labeled $DistanceToBatch$ that represents the time of loan application submission minus the assigned batch cutoff time.

For each loan j of applicant i, we run the following regression:

$$\begin{aligned} Y_{ij} = \quad & \beta_1 DistanceToBatch_{ij} + \beta_2 PostBatch_{ij} + \\ & \beta_3 DistanceToBatch_{ij} \times PostBatch_{ij} + \delta X_{ij} + \epsilon_{ij} \end{aligned} \tag{1}$$

where X controls for individual borrower characteristics and a variety of application time fixed effects (hour-of-day, day-of-week, and month). Our main outcome variable is whether the loan was repaid. The coefficient β_2 identifies the effect of missing a batch cutoff under the assumption that borrowers near the cutoff (on either side) are similar in terms of ex ante repayment/default likelihood.

To estimate Equation (1) and plot the results, we use the `rdrobust` suite of commands developed by Calonico, Cattaneo, Farrell, and Titiunik (2017). The commands allow for optimal bandwidth selection and automatically provide confidence intervals robust to bias induced by the optimal bandwidth selection. We also report specifications with fixed two-hour bandwidths, that exactly match our discontinuity figures. Because $PostBatch$ is assigned at the loan level, we do not cluster standard errors.[8]

3.1 Data construction

Our empirical strategy requires the identification of batches and batch times. Here we outline our procedure and refer to Appendix B for additional details.

Constructing batches We do not explicitly observe the batch to which a loan is assigned, nor we know when a batch is submitted to the bank for disbursement. In our example shown in Figure 1, this means that we do not observe the batches' disbursement times t_2 and t_4. In any given day, most loan deposit times are bunched together in time, and within a bunch, they are disbursed within seconds or milliseconds from one another. Therefore, we use a K-means clustering algorithm for disbursement times to reconstruct the batches for each day.

Constructing the cutoffs Next, for each batch, we determine the batch cutoff times (e.g. t_1 and t_3 in Figure 1). Recall that these times represent the latest moment a loan

[8] See Abadie, Athey, Imbens, and Wooldridge (2017).

could have been received and processed in the existing batch. Since batch cutoff times are not observable, we use the submission time of the last loan that is included in the batch as a proxy. For example, in Figure 1, our proxy for t_1 is given by the application submission time for loan j. We refer to these cutoffs as lower-bound cutoffs (LBCs hereafter), as they precede the actual, unobserved cutoff t_1. The loan that generates the LBC is labeled the LBC loan.[9] Finally, we assign each loan in the sample to the closest LBC, and code $DistanceToBatch$ and $PostBatch$ accordingly.[10] As mentioned earlier, loans for repeat borrowers are processed more quickly than loans for first-time borrowers. Thus, we calculate separate LBCs for first-time and repeat borrowers and run this procedure separately for the two types of loans.

Figure 2: **Impact of the cutoff on the likelihood of loan processing in batches**

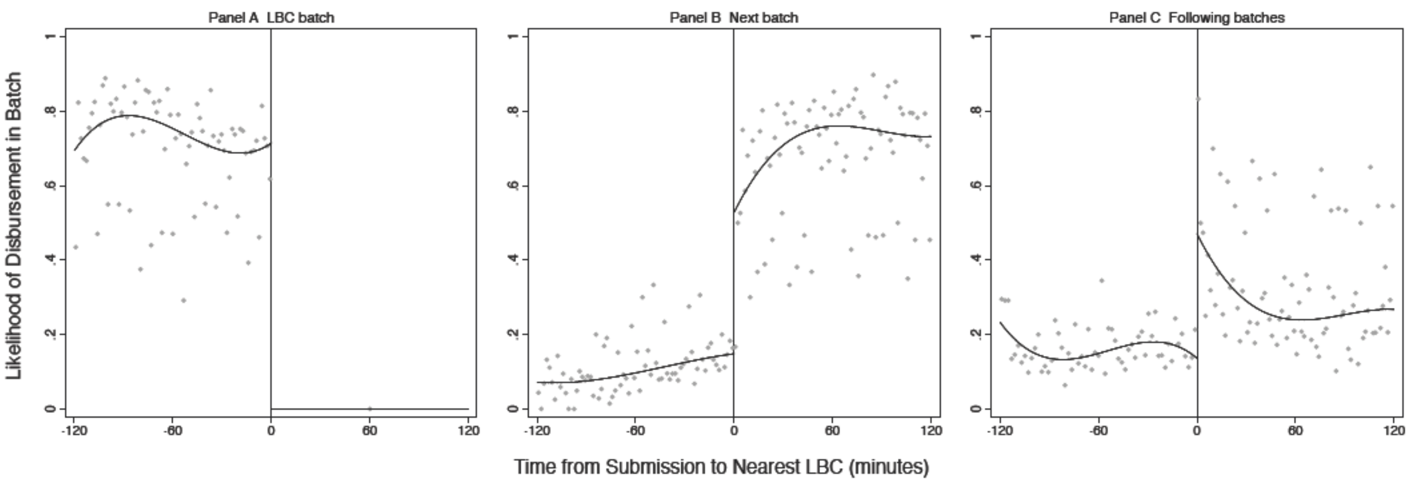

Notes: Regression discontinuity plots of the likelihood of disbursement in the LBC batch, the next batch, or the following batches. The RD uses third degree polynomials, a uniform kernel, and a fixed bandwidth of 120 minutes. The vertical line at 0 refers to the LBC loan in the batch considered in Panel A. We exclude the LBC loan.
By construction, in Panel A there are no observations after the LBC cutoff, as all loans after the LBC loan are processed in future batches. Loans submitted prior to the LBC cutoff can appear on the left hand sides of Panels A, B, and C, depending on the length of the verification process. Loans submitted after the LBC cutoff can appear only on the right hand sides of Panels B and C.

Figure 2 shows the results of our procedure by plotting the likelihood that a loan is

[9]To be clear, the reason this procedure yields a lower bound of the batch cutoff is because we cannot know whether any loan received between the LBC loan and the next observed loan could have been in the same batch as the LBC loan or if it would have been in a subsequent batch.

[10]Recall that in any given day, there are multiple batches, and therefore multiple batch cutoffs. To use each loan as a single observation, some assignment rule is necessary.

processed in the same batch as the LBC loan (Panel A), in the next batch (Panel B), or in the following batches (Panel C) as a function of *DistanceToBatch* and the LBC (which is centered at zero). In total, 70% of the loans issued before the cutoff are disbursed within the same batch as the LBC loan. Because of the way the LBC is constructed, there are no loans after the LBC time (Panel A) in the LBC batch. Panel B and, to a lesser extent, Panel C show that the likelihood of a loan being processed in subsequent batches jumps immediately after the LBC. The discontinuity is very sharp for repeat loans and less clearly defined for first-time loans (see Appendix Figures A1 and A2). This finding is in line with the expectation that there is more volatility in the length of time it takes to verify a first-time borrower than a repeat borrower.

3.2 Cutoffs and selection

Finally, we discuss three issues that arise with our approach and their solution. First, the density of submission times after an LBC is lower than the density before it (see Appendix Figure A3). This is not due to active manipulation by the applicants or the lender; it arises mechanically since by definition, an LBC is the submission time of the last loan that is included in the batch.[11]

Second, LBC loans are different from other loans in that they are processed quickly. The average delay in disbursing LBC loans is 4.4 hours, compared to 10.7 hours for applications submitted in the five minutes prior to the LBC. One reason for this is that LBC loans are selected on speed; if a loan is unable to be processed quickly, it is unlikely to become the last loan to make it into a batch. Thus, it is likely that LBC loans might also be different along unobservables.

Third, loans submitted just after the LBC may be more difficult to process than loans

[11] This phenomenon is similar to what Miller and Sanjurjo (2018) call "streak selection bias" in the context of collecting data to analyze the hot hand fallacy, and can be shown in a simulation of our data with a uniform density of submission times.

submitted just before it. If they were not, they would have been included in the batch with the LBC loan and become the LBC themselves. Figure 2 provides a visual confirmation that loan applications submitted shortly after the LBC (within approximately the next 20 minutes) are different from later applications: they have a lower likelihood of being processed in the next batch (Panel B) and are more likely to be processed in future batches (Panel C). Appendix B.3 provides additional evidence that loan applications submitted within 20 minutes after the LBC are negatively selected along observables.

We address these issues by dropping from the analysis LBC loans and all loans received within 20 minutes after the LBC (863 and 512 loans, respectively). In other words, we employ a one-sided "half-doughnut RD," where we drop only the right side of the doughnut hole.[12] This process yields our estimation sample of 11,512 loans. With the half-doughnut RD, the loan submission density and the observable characteristics of the borrower are smooth across the cutoff (see evidence in Appendix B.3). This process has the added benefit of reducing measurement error associated with proxying for the cutoff as long as we do not overshoot the true cutoff by more than its distance to the LBC. As shown in the following Section our results do not vary with the size of the half doughnut (i.e. post-LBC exclusion window).

4 Results

4.1 First stage

We begin by showing that the batching process causes loan applications submitted after LBCs to be disbursed with longer delays. To do so, we estimate Equation (1) using the delay length (in hours) as the dependent variable. We winsorize the delay distribution at the 90th percentile to account for a large right tail that is not of interest: the longest delay

[12]Note that selection concerns are absent for the loans that were submitted before the LBC because those loans are processed in either the same batch as the LBC or in following batches, i.e., they are not selected based on their batching. We thus include all loans leading up to the LBC.

in our sample is over 27 days, while the 90th percentile delay is 63 hours.

Figure 3, Panel A, displays the half-doughnut RD plots for loan delays. In these plots, we assume a bandwidth of two hours around the LBC; we estimate a linear fit; and we use a uniform estimation kernel. There is a clear increase in delays at the LBC. The relative size of this effect is more pronounced for repeat loans than for first-time loans. This is because, as mentioned earlier, the average delivery speed of repeat loans is higher than that of first-time loans.

Figure 3: **Half-doughnut RD plots**

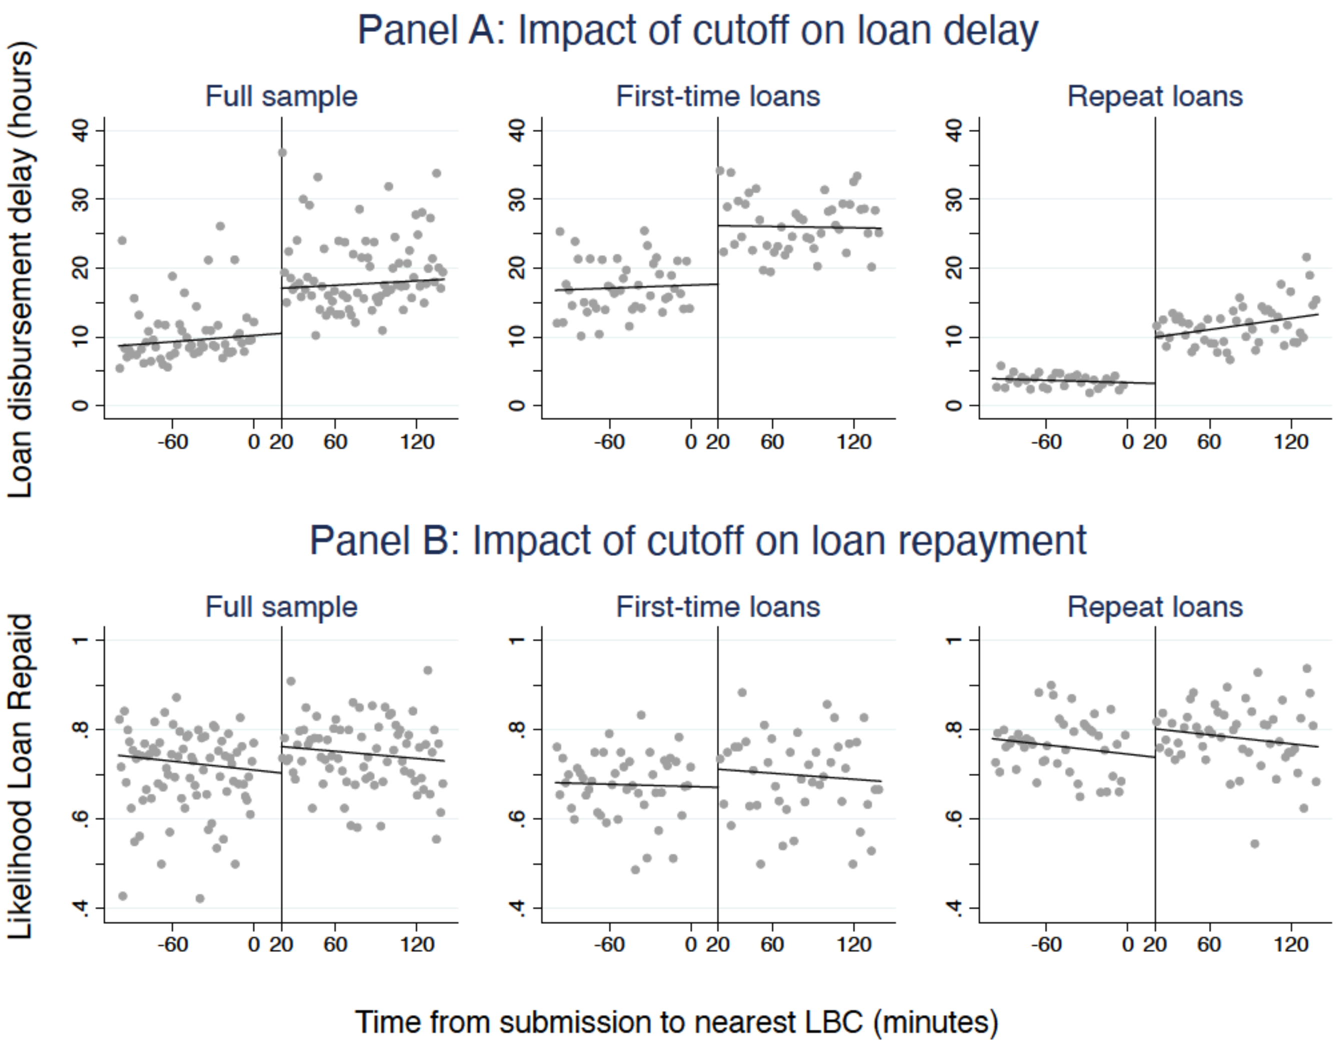

Notes: Regression discontinuity plots use linear fit with a uniform kernel and a fixed bandwidth of 120 minutes. LBC loans and all loans received within 20 minutes after the LBC.

Appendix Table A3 reports RD estimates using both a model that exactly matches the specification in Figure 3 and optimal-bandwidth models controlling for borrower demographics and application submission time fixed effects. In every specification, there is a large and statistically significant effect of the cutoff on loan delay. We estimate that missing the cutoff increases the borrower's wait time by almost 10 hours, effectively doubling the wait time. The increase in the delay is similar for first-time and repeat loans (11 and 8 hours, respectively), which implies a 63% increase in delays for first-time loans and a 228% increase for repeat loans. In addition, the induced delays greatly decreases the likelihood of same-day disbursement. Appendix Table A4 shows that the impact of missing a batch cutoff on the likelihood that a borrower receives her loan on the same day falls by 24 percentage points.

4.2 Main results: effect of delays on repayment

We now estimate the effects of the delay-inducing cutoff on loan repayment rates. Figure 3, Panel B, displays the half-doughnut RD plots for loan repayment. We observe an increase in the likelihood of repayment at the 20-minute post-LBC cutoff for the full sample, first-time loans, and repeat loans. The corresponding regression estimates are reported in Table 1. The specification in column (1) matches Figure 3: it uses a two-hour bandwidth, uniform kernel, and linear estimation. Columns (2)-(4) use optimal bandwidth selection and a triangular estimation kernel. We allow for an asymmetric optimal bandwidth because the exclusion of loans submitted within 20 minutes following the LBC creates an asymmetry in density around the post-LBC latent cutoff. Panel A shows the full sample estimates, and Panels B and C show estimates for first-time and repeat loans, respectively. Below each estimate, we report the following information: the heteroskedasticity-robust p-values of the linear estimates; the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates;[13] the effect magnitude as a percentage of the pre-cutoff mean repayment within

[13]The first p-value has the advantage of pertaining to the point estimate of interest, but it does not account for potential bias due to bandwidth selection. The second one accounts for bias due to bandwidth

two hours of the cutoff; the optimal bandwidth as determined by the rdrobust command; and the number of observations within that optimal bandwidth.[14]

For the full sample, the induced delay (10 hours on average) increases repayment rates by six percentage points, which corresponds to an 8% increase in repayment rates (equivalently, a 21% reduction in the default rate). The effect is similar in magnitude across specifications and is always statistically significant according to both sets of p-values. Column (4) shows a statistically significant 7.4 percentage point (10%) increase in repayment for repeat loans and an almost statistically significant 5.4 percentage point (8%) increase in repayment for first-time loans. The differences in estimates however, are not statistically significant.

These results demonstrate a causal effect of induced delays on repayment: a 5.6 percentage point increase in repayment in response to an induced additional delay of 9.81 hours (estimates from column (4) of Panel A in Table 1 and Appendix Table A3, respectively). Back-of-the-envelope calculations imply an increase of 0.6 percentage points per hour of induced additional delay. Alternatively, we can directly estimate the causal effect of loan disbursement delay on repayment rates (albeit at the same margin as the crude calculation) using two-stage least squares. We instrument for loan disbursement delay using our regression discontinuity model from Equation (1) using a fixed bandwidth of two hours.[15] This approach yields slightly smaller but qualitatively similar results; using the most robust specification in the full sample, we estimate that each hour of induced delay increases repayment rates by 0.4 percentage points ($p = 0.016$). Estimates are shown in Appendix Table A5. Finally, Appendix Figure A4 shows that the estimates are robust to the post-LBC exclusion window.

selection, but it pertains to the quadratic estimate used for bias correction, not the linear estimate of interest.

[14]The sample that is fed into the optimal bandwidth algorithm is held fixed across specifications. The number of observations within the optimal bandwidth varies slightly across specifications as the optimal bandwidth changes when adding controls.

[15]The use of least squares implies a uniform estimation kernel.

Table 1: **Impact of the cutoff on loan repayment**

RD bandwidth:	Two-hour	Optimal		
	(1)	(2)	(3)	(4)
A. Full sample (N = 11,512)				
PostBatch	0.059	0.063	0.060	0.056
	(0.023)	(0.024)	(0.024)	(0.024)
Estimate p-value	0.011	0.008	0.012	0.018
Bias-corrected estimate p-value	0.044	0.017	0.021	0.038
Effect as % of pre-cutoff mean	8%	9%	8%	8%
Optimal bandwidth (mins)		[144,119]	[144,112]	[146,112]
Observations within bandwidth	7,177	7,704	7,602	7,658
B. First-time loans (N = 5,530)				
PostBatch	0.041	0.040	0.041	0.054
	(0.036)	(0.035)	(0.034)	(0.034)
Estimate p-value	0.259	0.251	0.227	0.110
Bias-corrected estimate p-value	0.813	0.326	0.274	0.146
Effect as % of pre-cutoff mean	6%	6%	6%	8%
Optimal bandwidth (mins)		[153,136]	[162,126]	[164,126]
Observations within bandwidth	3,090	3,565	3,554	3,577
C. Repeat loans (N = 5,982)				
PostBatch	0.064	0.083	0.078	0.074
	(0.030)	(0.034)	(0.034)	(0.034)
Estimate p-value	0.037	0.015	0.021	0.029
Bias-corrected estimate p-value	0.015	0.038	0.050	0.067
Effect as % of pre-cutoff mean	8%	11%	10%	10%
Optimal bandwidth (mins)		[123,110]	[127,110]	[123,111]
Observations within bandwidth	4,087	4,036	4,084	4,068
Day-of-week, hour-of-day, month FEs	N	N	Y	Y
Borrower controls	N	N	N	Y

Notes: Estimates exclude LBC loan, and loans received within 20 minutes after the LBC. Column (1) reports a specification with a uniform estimation kernel and a fixed bandwidth of 120 minutes around the 20-minute post-LBC cutoff. Columns (2-4) report specifications with a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. Heteroskedasticity-robust standard errors of the linear estimates are shown in parentheses below the estimates, calculated using the nearest-neighbor variance estimator with a minimum of three matches. Below each estimate, we report: the heteroskedasticity-robust p-values of the linear estimates; the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates; the estimated effect as a percentage of the pre-cutoff mean repayment rate within the two-hour bandwidth; the optimal bandwidths, rounded to the nearest integer (for the specifications in columns (2)-(4)); and observations within the used bandwidth. The overall sample sizes for each panel correspond to all loans within twelve hours of an LBC. Column (2) has no control variables, column (3) controls for application submission day-of-week, hour-of-day and month fixed effects, and column (4) adds borrower controls (age, age squared, sex, marital status, number of dependents, log income, and credit score). In Panels A and C, we also add a fixed effect for a borrower's sequential loan number in column (3).

4.3 Heterogeneity analysis

In Appendix Table A6, we report the impact of induced delays on repayment by marital status (single/divorced/widowed vs. married), income (below/above median), and creditworthiness. We find an effect of 10.8 percentage points ($p = 0.003$) for the married sample and null effects for the unmarried sample. Regarding income, we find an effect of 7.9 percentage points ($p = 0.011$) for individuals with above-median income and an effect of 2.8 percentage points ($p = 0.424$) for individuals with below-median income. Additional estimates show an effect of 14.2 percentage points ($p = 0.001$) for borrowers assessed by the lender to have a "better" or "best" credit score and of 3.3 percentage points ($p = 0.251$) for those rated "average," "marginal," or "none."

In addition, Appendix Table A7 splits the sample between applications received before and after noon. Afternoon applications remain longer in batches and are more likely to be delayed until the next day than morning applications. The effects are stronger for the former group (7 percentage points, $p = 0.02$) relative to morning applications (2.8 percentage points, $p = 0.53$).

4.4 Timing of repayment

Our analysis thus far has considered the effect of delays on whether a loan was repaid. Now, we study when loans are repaid. For this analysis, we rearrange our data as a panel. For each loan in the sample, we define the time dimension as the number of days since the loan was disbursed, ranging from zero to 356 (the latest repayment we observe). For each loan observation day, a loan is classified as repaid or not. We estimate the effect of missing a batch cutoff using our regression discontinuity specification one day at a time. These estimates measure the difference in repayments for each date.

Figure 4 plots the RD estimates over time. Panel A reports the estimates for the repayment periods 1-30 days after loan disbursement. We label the period before seven days as the "early" repayment period because the shortest possible loan term is seven days (note

Figure 4: **RD estimates over time**

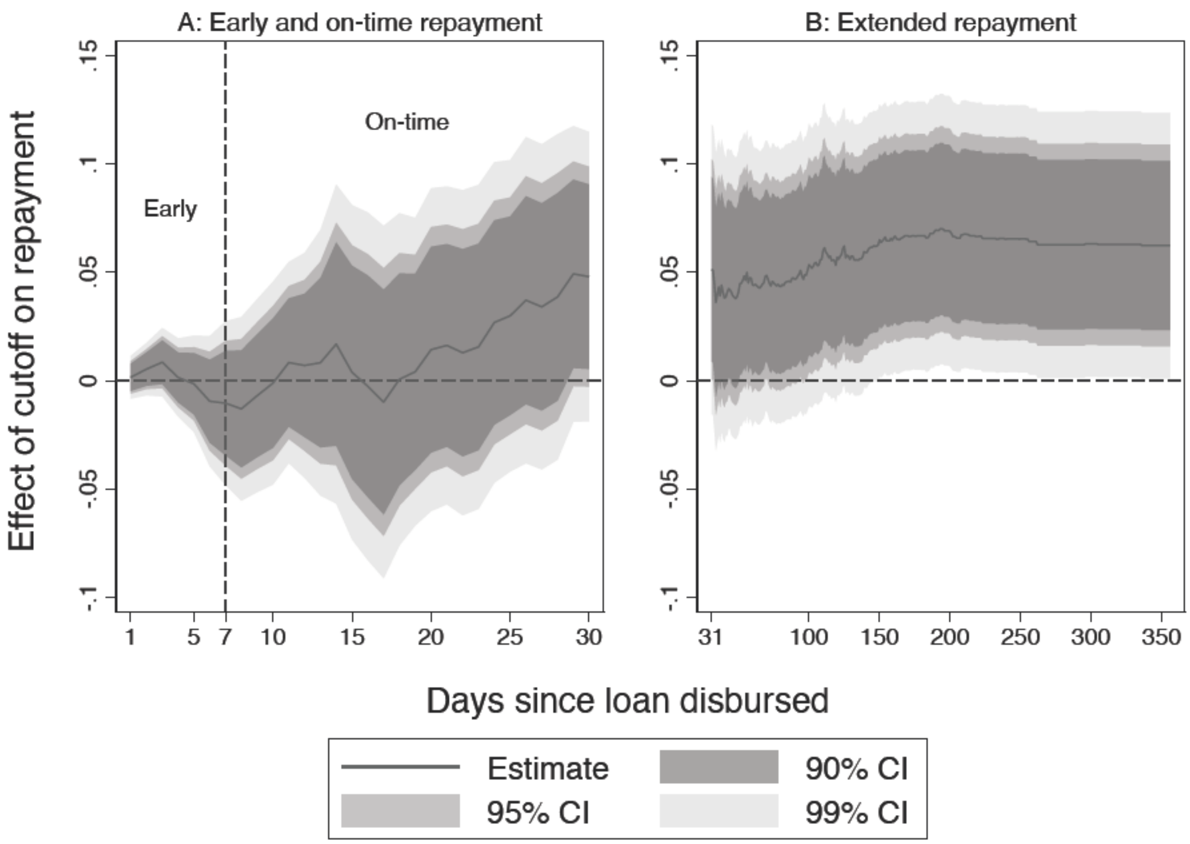

Notes: RD estimates plots, using specification from column (2) of Table 1 on whether loan was paid a certain number of days after the issuance of the loan. We use the conventional confidence intervals for the figure because they pertain to the estimated coefficient.

that we do not directly observe the contracted term). Panel A clearly shows that there is no difference in the repayment behavior of delayed loans in the "early" repayment period. Differences in repayment begin to emerge only 17 days after disbursement. Panel B reports the estimates for the repayment period 30-365 days after loan disbursement. After 30 days, we can begin to detect a significant effect of the cutoff. We estimate an effect of the cutoff of 4.8 percentage points ($p = 0.064$, bias-corrected $p = 0.082$), which represents roughly three-quarters of the overall effect, on day 30. During the extended repayment period, the slope remains positive, explaining the remaining effect. The point estimate at the end of the extended repayment period corresponds to the RD estimate in Table 1.[16] As a final check, we separately analyze loans submitted between November and January from later loans. Estimates are similar in both samples, confirming that additional time to repay a loan does not have an effect on repayments.

4.5 Lender's profitability

Finally, we study the effects of loan delays on future borrowing behavior and on the profits of the lender. First-time borrowers whose loans are delayed may be more likely to repay their initial loans and, consequently, may be more likely to borrow as they become eligible to borrow again. At the same time, they might reduce their demand for credit if they believe the lender is too "slow". In Appendix C, we show that there are positive but statistically insignificant delays on the likelihood of borrowing again, on the repayment behavior of future loans, and on the total number of loans taken. Lacking evidence of negative effects of the delay, we argue that the overall impact of delays on the profitability of the lender is positive.[17]

4.6 Mechanisms

Several mechanisms might explain our findings. Despite the limited administrative data, we are able to speculate about the likely channels and exclude others.

Loan declines and early repayments We first rule out the possibility that borrowers facing a delay decline the loan before it is issued. This could explain our findings if loan declines are disproportionally found among borrowers with a low likelihood of repayment. We obtained from the lender a separate dataset of successful applications that ended in the

[16] We do not observe the loan term for repeat loans. Hence, we do not carry out an analysis of timely loan repayments. When we break down first-time loans by their duration, we obtain point estimates that are consistent with the findings in this section, but are also noisy.

[17] Note, however, that as we do not have information on the cost side of the firm, our welfare analysis is limited in its ability to quantify the effects of the intervention on firm profits.

client rejecting the loan prior to disbursement. For the study period, we identified a total of 557 approved loans that were rejected by the applicant prior to disbursement. These make up 2.5% of all loans disbursed, a fraction too small to drive the results.

A second possibility is that clients returned delayed loans immediately after disbursement. However, Figure 4 shows that there is no difference in repayments between delayed and immediate loans disbursed in that time period. Moreover, only 8% of all loans were returned within seven days of disbursement.

Increased deliberation A plausible explanation for our results is that disbursement delays provide borrowers with extra time to deliberate about the use of their approved loans. Existing research suggests that waiting periods (which provide the time for deliberation) improve the consumption choices individuals make (Imas et al., 2016; DeJarnette, 2018; Brownback et al., 2019),[18] and could induce borrowers to make a repayment plan (Thakral and Tô, 2020). In our context, increased deliberation could convince borrowers to change the use of the loan so that they have more liquidity at the time of repayment.[19] Alternatively, it might induce borrowers to develop a robust repayment plan. Unfortunately, our administrative data do not contain information about the intended or actual use of loans or about borrowers' repayment plans.

Household dynamics As discussed earlier, the effect of delays is stronger for married applicants and for applications submitted in the afternoon, which are more likely to be delayed overnight. We speculate that without a delay, an individual may be able to apply for, obtain, and use a loan without confronting their partners, while household bargaining

[18]Imas et al. (2016) find that enforcing waiting periods to temporally separate the news about a new consumption choice set from the ability to make a choice from that set leads to a substantial increase in patient choices.

[19]This situation presumes a certain elasticity in the use of the loan. Evidence from the microfinance industry suggests that credit use is flexible and responds to the characteristics of the loan (Field et al., 2013).

becomes an issue if disbursement is delayed overnight. Intrahousehold negotiations could improve repayments through deliberation (as discussed above) or through a pooling of resources. Further analysis in Appendix Table A8 indicates that the effect of marital status is mediated by gender. The effect of the delay for married women is 18.3 percentage points ($p = 0.002$), while for unmarried women, it is -3.1 percentage points ($p = 0.464$), while there are no statistically significant differences for married and unmarried men. These effects point to potentially interesting intrahousehold dynamics that merit further study but are beyond the scope of this paper due to space and data limitations.

Time-sensitive loan needs Borrowers facing time-sensitive consumption or investment opportunities that expire before loans are delivered might not want a loan after it is received. Higher repayments could be explained by the fact that funds have been unused. Alternatively, delayed borrowers with urgent needs could seek alternative sources of credit from other digital lenders. This additional credit could provide the necessary liquidity to repay delayed loans but at the cost of a higher level of overall debt.[20]

5 Conclusion

We study whether one of the primary features of digital credit—the speed of delivery of funds—affects the likelihood that a loan is repaid. To date, despite the continuous growth of this market, this question remains unanswered. This is partly because detailed administrative data are not easily available. Our study combines difficult-to-obtain administrative data from a digital lender with a robust identification strategy, and shows that reducing the speed of delivery of digital loans increases the likelihood that loans are repaid by 6 percentage points.

[20]While we cannot directly explore credit use with our data, we can explore the role of liquidity on repayments. We are able to rule out the earnings cycle as a confounding factor: we replicate our results after dropping loans that are due within two days of payday (mid-month and end of month) and our results remain very similar.

This corresponds to a 21% reduction in the likelihood of loan default.

These findings naturally raise the question of whether regulating the speed of digital credit disbursement, such as by imposing a waiting period on loan delivery, could protect consumers from avoidable defaults. While our analysis is suggestive, the full answer requires a careful welfare analysis. In our setting, a number of mechanisms are consistent with our results, so it is unclear whether the overall effect of delays on borrowers is positive. On the one hand, higher repayments lead to higher credit scores and improved future loan terms. However, we cannot rule out the possibility that consumers miss out on timely and profitable opportunities, are unable to address an immediate need, or address their need by taking loans from other sources and increasing their overall indebtedness. We can be more conclusive about the effect of delays to the lender: profits are higher for delayed loans. Overall, our study justifies further work on mandatory waiting periods for digital credit as a potential consumer protection measure.

References

Abadie, A., S. Athey, G. W. Imbens, and J. Wooldridge (2017). When should you adjust standard errors for clustering? Technical report, National Bureau of Economic Research.

Ashraf, N., D. Karlan, and W. Yin (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics 121*(2), 635–672.

Bauer, M., J. Chytilová, and J. Morduch (2012). Behavioral foundations of microcredit: Experimental and survey evidence from rural india. *American Economic Review 102*(2), 1118–39.

Bharadwaj, P., W. Jack, and T. Suri (2019). Fintech and household resilience to shocks: Evidence from digital loans in kenya. Technical report, National Bureau of Economic Research.

Björkegren, D. and D. Grissen (2018). The potential of digital credit to bank the poor. In *AEA Papers and Proceedings*, Volume 108, pp. 68–71.

Brooks, A. W. (2015). Emotion and the art of negotiation. *Harvard Business Review*.

Brownback, A., A. Imas, and M. A. Kuhn (2019). Behavioral interventions increase the effectiveness of healthy food subsidies.

Burgess, H. (2004). *Beyond Intractability*, Chapter Cooling-Off Periods. Conflict Information Consortium, University of Colorado, Boulder.

Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2017). rdrobust: Software for regression-discontinuity designs. *The Stata Journal 17*(2), 372–404.

Carlson, S. (2017). Dynamic incentives in credit markets: An exploration of repayment decisions on digital credit in africa. *MIT, Cambridge, MA, USA Department of Economics*.

Cattaneo, M. D., M. Jansson, and X. Ma (2018). Manipulation testing based on density discontinuity. *The Stata Journal 18*(1), 234–261.

DeJarnette, P. (2018). Temptation over time: Delays help. Technical report, Working paper.

Donovan, K. P. and E. Park (2019). Perpetual debt in the silicon savannah.

Duflo, E., M. Kremer, and J. Robinson (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American economic review 101*(6), 2350–90.

Dupas, P. and J. Robinson (2013). Why don't the poor save more? evidence from health savings experiments. *American Economic Review 103*(4), 1138–71.

Edwards, G., E. Nesson, J. J. Robinson, and F. Vars (2018). Looking down the barrel of a loaded gun: The effect of mandatory handgun purchase delays on homicide and suicide. *The Economic Journal 128*(616), 3117–3140.

Feigenberg, B., E. Field, and R. Pande (2013). The economic returns to social interaction: Experimental evidence from microfinance. *Review of Economic Studies 80*(4), 1459–1483.

Field, E., R. Pande, J. Papp, and N. Rigol (2013). Does the classic microfinance model discourage entrepreneurship among the poor? experimental evidence from india. *American Economic Review 103*(6), 2196–2226.

Francis, E., J. Blumenstock, and J. Robinson (2017). Digital credit: A snapshot of the current landscape and open research questions. *CEGA White Paper*.

Heidhues, P. and B. Kőszegi (2010). Exploiting naivete about self-control in the credit market. *American Economic Review 100*(5), 2279–2303.

Imas, A., M. Kuhn, and V. Mironova (2016). Waiting to choose.

Karlan, D., M. Morten, and J. Zinman (2015). A personal touch: Text messaging for loan repayment. *Behavioral Science and Policy 1*(2), 25–31.

Karlan, D. and J. Zinman (2009). Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica 77*(6), 1993–2008.

Karlan, D. and J. Zinman (2010). Expanding credit access: Using randomized supply decisions to estimate the impacts. *The Review of Financial Studies 23*(1), 433–464.

Koenig, C. and D. Schindler (2018). Dynamics in gun ownership and crime-evidence from the aftermath of sandy hook. Technical report, Working paper.

Kremer, M., J. Lee, J. Robinson, and O. Rostapshova (2013). Behavioral biases and firm behavior: Evidence from kenyan retail shops. *American Economic Review 103*(3), 362–68.

Makles, A. (2012). Stata tip 110: How to get the optimal k-means cluster solution. *The Stata Journal 12*(2), 347–351.

Malingha, D. (2019, August). This nobel prize-winning idea is instead piling debt on millions. Bloomberg Future Finance.

Mazer, R. and A. Fiorillo (2015). Digital credit: Consumer protection for m-shwari and m-pawa users. *CGAP. April 21*.

McKee, K., M. Kaffenberger, and J. M. Zimmerman (2015). Doing digital finance right: The case for stronger mitigation of customer risks. *CGAP Focus Note 103*.

Meier, S. and C. Sprenger (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics 2*(1), 193–210.

Melzer, B. T. (2011). The real costs of credit access: Evidence from the payday lending market. *The Quarterly Journal of Economics 126*(1), 517–555.

Miller, J. B. and A. Sanjurjo (2018). Surprised by the hot hand fallacy? a truth in the law of small numbers. *Econometrica 86*(6), 2019–2047.

Morse, A. (2011). Payday lenders: Heroes or villains? *Journal of Financial Economics 102*(1), 28–44.

Skiba, P. M. and J. Tobacman (2019). Do payday loans cause bankruptcy? *The Journal of Law and Economics 62*(3), 485–519.

Thakral, N. and L. T. Tô (2020). Anticipation and consumption. *Mimeo*.

A Appendix for online publication

Figure A1: **Impact of the cutoff on the likelihood of loan processing in batches (first-time loans)**

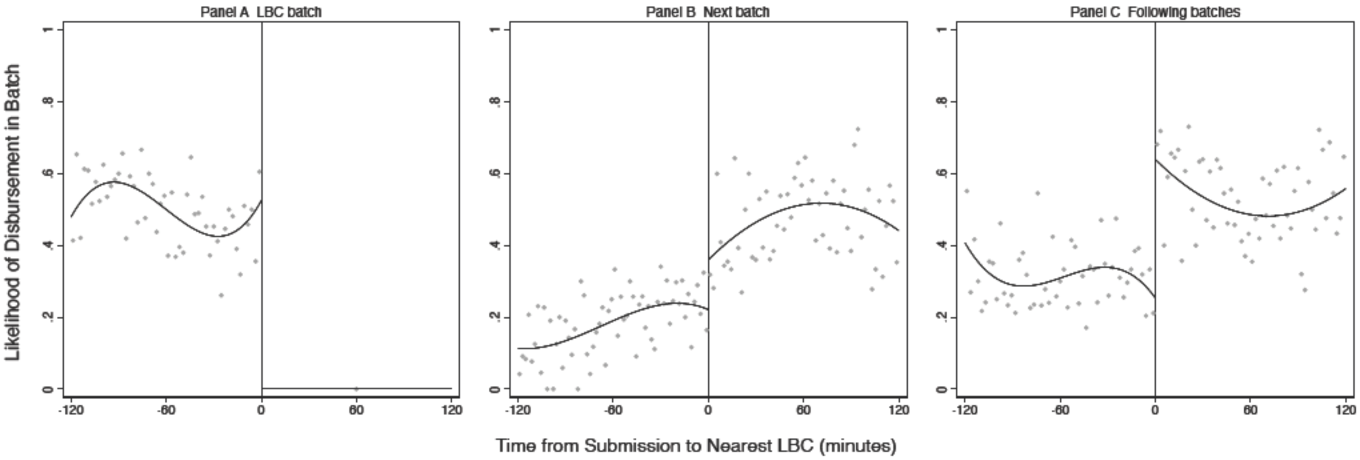

Notes: RD plots use third degree polynomials, a uniform kernel, and a fixed bandwidth of 120 minutes. LBC loans excluded.

Figure A2: **Impact of the cutoff on the likelihood of loan processing in batches (repeat loans)**

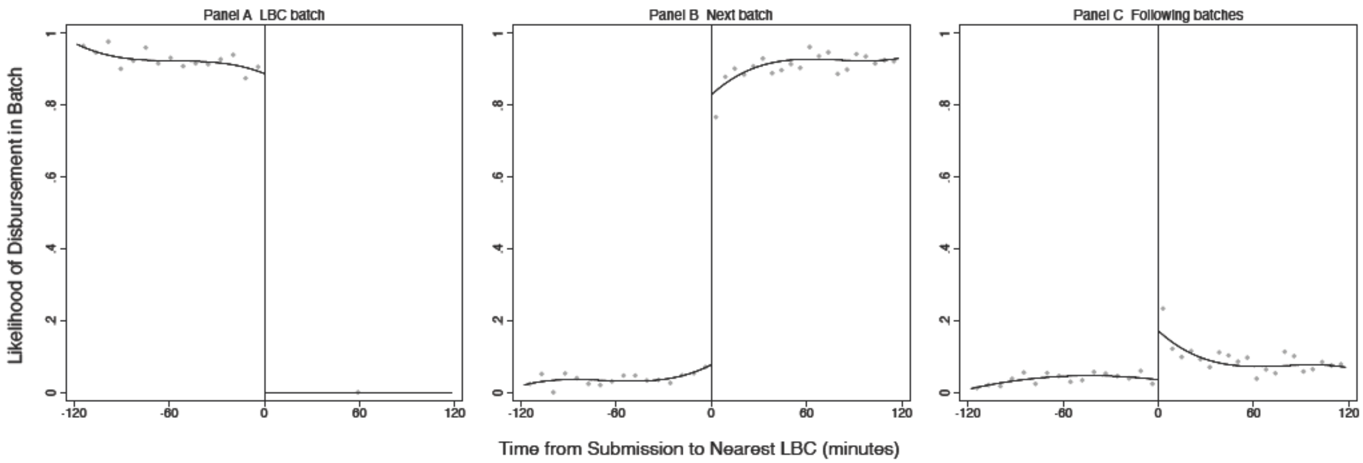

Notes: RD plots use third degree polynomials, a uniform kernel, and a fixed bandwidth of 120 minutes. LBC loans excluded.

Figure A3: **Density of** $DistanceToBatch$**, 12-hour window**

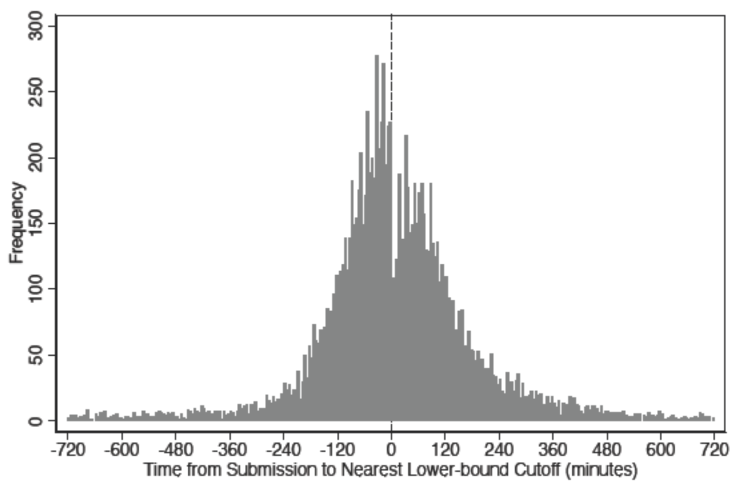

Notes: Five-minute bins.

Figure A4: **Impact of the cutoff on loan repayments by post-LBC cutoff**

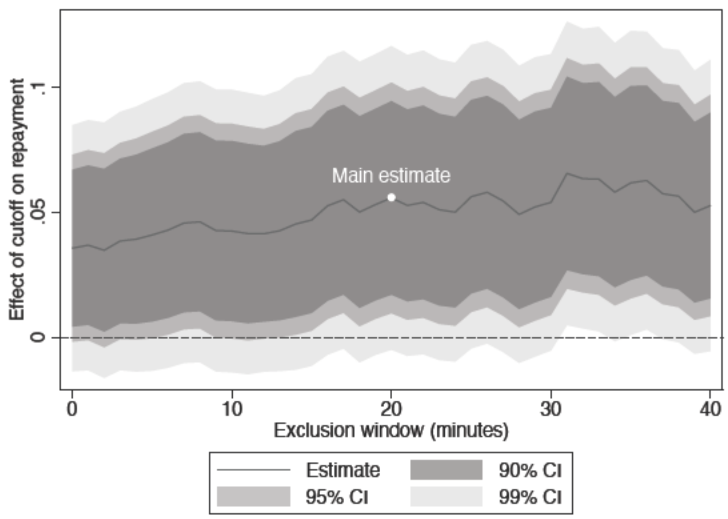

Notes: Estimates are from the same model as Table 1, column (4), estimated for each post-LBC exclusion window from zero to forty (in one-minute increments).

Table A1: **Summary statistics**

Variables	Mean	SD	Min	Median	Max
A. Borrower characteristics (N = 7,206)					
Age	37.45	9.55	20	36	65
Female	0.4	0.50	0	0	1
Married	0.49	0.50	0	0	1
Dependents	1.24	1.14	0	1	5
Monthly income (pesos)	1,718.66	8,279.59	291.67	916.67	125,000.00
Credit score - none	0.13	0.33	0	0	1
Credit score - marginal	0.30	0.46	0	0	1
Credit score - average	0.31	0.46	0	0	1
Credit score - better	0.22	0.41	0	0	1
Credit score - best	0.04	0.21	0	0	1
Credit score - linear (0-4)	1.76	1.07	0	2	4
B: All loans (N = 11,512)					
Delay (hours)	16.00	19.65	0.15	5.10	63.10
Loan repaid	0.73	0.44	0	1	1
C: First-time loans (N = 5,530)					
Amount received (pesos)	1,759.29	348.53	1,000	1,500	3,000
Loan term (days)	21.36	7.13	7	21	30
Delay (hours)	23.63	21.32	0.60	18.07	63.10
Loan repaid	0.68	0.46	0	1	1
D: Repeat loans (N = 5,982)					
Delay (hours)	8.95	14.84	0.15	2.99	63.10
Loan repaid	0.78	0.42	0	1	1

Notes: Borrower characteristics are collected at the time of the first loan application. Income is winsorized at the top 0.5% due to a couple extreme outliers. Loan amounts and lengths are only available for first loans. Delays measure the time between loan application and loan disbursement. Delays are winsorized at the top 10% due to a large right tail.

Table A2: **Borrower/loan characteristics and loan repayment**

Sample:	Full sample		First-time loans	
	(1)	(2)	(3)	(4)
Age	-0.004	-0.004	-0.008	-0.009
	(0.003)	(0.003)	(0.005)	(0.005)
	$p = 0.274$	$p = 0.218$	$p = 0.103$	$p = 0.081$
Age^2	0.000052	0.000055	0.000096	0.000103
	(0.000041)	(0.000040)	(0.000062)	(0.000062)
	$p = 0.211$	$p = 0.172$	$p = 0.123$	$p = 0.099$
Female	0.013	0.011	0.018	0.018
	(0.009)	(0.008)	(0.013)	(0.013)
	$p = 0.151$	$p = 0.177$	$p = 0.154$	$p = 0.154$
Married	-0.010	-0.011	-0.014	-0.015
	(0.010)	(0.010)	(0.015)	(0.015)
	$p = 0.292$	$p = 0.251$	$p = 0.339$	$p = 0.318$
Dependents	-0.007	-0.006	-0.002	-0.001
	(0.004)	(0.004)	(0.007)	(0.007)
	$p = 0.116$	$p = 0.142$	$p = 0.812$	$p = 0.937$
Log monthly income (pesos)	0.023	0.020	0.013	0.011
	(0.005)	(0.005)	(0.008)	(0.008)
	$p < 0.001$	$p < 0.001$	$p = 0.117$	$p = 0.172$
Credit score (0-4)	0.026	0.034	0.065	0.073
	(0.004)	(0.004)	(0.009)	(0.009)
	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
Log amount received (pesos)			0.029	0.003
			(0.043)	(0.045)
			$p = 0.502$	$p = 0.955$
Loan term (days)			-0.003	-0.003
			(0.001)	(0.001)
			$p = 0.001$	$p = 0.002$
Day-of-week, hour-of-day, month FEs	N	Y	N	Y
Observations	11,512	11,512	5,530	5,530
Clusters	7,206	7,206		
Sample mean [SD]	0.733 [0.442]		0.685 [0.465]	

Notes: All estimates are from linear probability models of repayment. Columns (1) and (2) use the entire estimation sample of loans, with standard errors clustered at the borrower level. Columns (3) and (4) use only first-time loans, with heteroskedasticity-robust standard errors. In columns (2) and (4), we include fixed effects for the hour-of-day, day-of-week, and month of application submission. In column (2) the set of fixed effects also includes a borrower's sequential loan number.

Table A3: **Impact of the cutoff on loan delay (in hours)**

RD bandwidth:	Two-hour	Optimal		
	(1)	(2)	(3)	(4)
A. Full sample (N = 11,512)				
PostBatch	6.56	10.76	9.85	9.81
	(0.87)	(1.50)	(1.25)	(1.08)
Effect as % of pre-cutoff mean	69%	113%	103%	103%
Optimal bandwidth (mins)		[81,49]	[95,53]	[132,55]
Observations within bandwidth	7,177	4,180	4,858	5,974
B. First-time loans (N = 5,530)				
PostBatch	8.50	12.25	11.18	10.91
	(1.57)	(2.00)	(1.77)	(1.76)
Effect as % of pre-cutoff mean	49%	71%	65%	63%
Optimal bandwidth (mins)		[129,62]	[122,72]	[123,72]
Observations within bandwidth	3,090	2,626	2,683	2,695
C. Repeat loans (N = 5,982)				
PostBatch	6.71	8.60	8.34	8.25
	(0.89)	(1.27)	(1.09)	(1.09)
Effect as % of pre-cutoff mean	185%	237%	230%	228%
Optimal bandwidth (mins)		[107,66]	[118,71]	[117,71]
Observations within bandwidth	4,087	3,189	3,426	3,426
Day-of-week, hour-of-day, month FEs	N	N	Y	Y
Borrower controls	N	N	N	Y

Notes: All estimated discontinuities are from linear models that exclude the LBC loan, and loans received within 20 minutes after the LBC. Dependent variable is the delay in disbursement. Column (1) reports a specification with a uniform estimation kernel and a fixed bandwidth of 120 minutes around the 20-minute post-LBC cutoff. Columns (2-4) report specifications with a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. Heteroskedasticity-robust standard errors of the linear estimates are shown in parentheses below the estimates, calculated using the nearest-neighbor variance estimator with a minimum of three matches. All estimates are statistically significant with $p < 0.001$ according to both the heteroskedasticity-robust p-values of the linear estimates, and the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates. We also report the estimated effect as a percentage of the pre-cutoff mean delay within the two-hour bandwidth. The optimal bandwidths –rounded to the nearest integer– are reported for the specifications in columns (2)-(4), and observations within the used bandwidth are reported below. The overall sample sizes for each panel correspond to all loans within twelve hours of an LBC. The fixed effects added in column (3) include the hour-of-day, day-of-week, and month of application submission. In Panels A and C, a fixed effect for the borrower's sequential loan number is also included. The borrower controls added in column (4) are age, age squared, sex, marital status, number of dependents, log income, and credit score.

Table A4: **Impact of the cutoff on the likelihood of same-day loan disbursement**

RD bandwidth:	Two-hour	Optimal		
	(1)	(2)	(3)	(4)
A. Full sample (N = 11,512)				
PostBatch	-0.148	-0.212	-0.231	-0.237
	(0.022)	(0.043)	(0.034)	(0.028)
Effect as % of pre-cutoff mean	-19%	-27%	-30%	-31%
Optimal bandwidth (mins)		[66,48]	[73,53]	[94,55]
Observations within bandwidth	7,177	3,582	4,097	4,873
B. First-time loans (N = 5,530)				
PostBatch	-0.191	-0.284	-0.264	-0.261
	(0.037)	(0.049)	(0.040)	(0.040)
Effect as % of pre-cutoff mean	-34%	-50%	-47%	-46%
Optimal bandwidth (mins)		[112,55]	[125,61]	[127,61]
Observations within bandwidth	3,090	2,371	2,573	2,592
C. Repeat loans (N = 5,982)				
PostBatch	-0.161	-0.179	-0.201	-0.198
	(0.025)	(0.035)	(0.025)	(0.025)
Effect as % of pre-cutoff mean	-17%	-19%	-21%	-21%
Optimal bandwidth (mins)		[95,74]	[108,89]	[117,71]
Observations within bandwidth	4,087	3,085	3,548	3,509
Day-of-week, hour-of-day, month FEs	N	N	Y	Y
Borrower controls	N	N	N	Y

Notes: All estimated discontinuities are from linear models that exclude the LBC loan, and loans received within 20 minutes after the LBC. Column (1) reports a specification with a uniform estimation kernel and a fixed bandwidth of 120 minutes around the 20-minute post-LBC cutoff. Columns (2-4) report specifications with a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. Heteroskedasticity-robust standard errors of the linear estimates are shown in parentheses below the estimates, calculated using the nearest-neighbor variance estimator with a minimum of three matches. All estimates are statistically significant with $p \leq 0.001$ according to both the heteroskedasticity-robust p-values of the linear estimates, and the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates. We also report the estimated effect as a percentage of the pre-cutoff mean likelihood of same-day disbursement within the two-hour bandwidth. The optimal bandwidths –rounded to the nearest integer– are reported for the specifications in columns (2)-(4), and observations within the used bandwidth are reported below. The overall sample sizes for each panel correspond to all loans within twelve hours of an LBC. The fixed effects added in column (3) include the hour-of-day, day-of-week, and month of application submission In Panels A and C, a fixed effect for the borrower's sequential loan number is also included. The borrower controls added in column (4) are age, age squared, sex, marital status, number of dependents, log income, and credit score.

Table A5: **IV estimates of the impact of loan delay on loan repayment**

	(1)	(2)	(3)
A. Full sample (N = 7,177)			
Loan delay (hours)	0.0026	0.0043	0.0042
	(0.0013)	(0.0018)	(0.0017)
Estimate p-value	0.041	0.014	0.016
B. First-time loans (N = 3,090)			
Loan delay (hours)	0.0026	0.0035	0.0042
	(0.0019)	(0.0027)	(0.0027)
Estimate p-value	0.172	0.193	0.123
C. Repeat loans (N = 4,087)			
Loan delay (hours)	0.0025	0.0048	0.0046
	(0.0017)	(0.0022)	(0.0022)
Estimate p-value	0.127	0.031	0.038
Day-of-week, hour-of-day, month FEs	N	Y	Y
Borrower controls	N	N	Y

Notes: All estimates are from two-stage-least-squares models where the regression-discontinuity specification from equation 1 instruments for the experienced delay in receiving a loan (jn hours). The sample limited to a two-hour window around the 20-minute post-LBC cutoff. Heteroskedasticity-robust standard errors are shown in parentheses below the estimates. All models feature first stages with joint F-statistics that are statistically different from zero with $p < 0.001$. The fixed effects added in column (2) include the hour-of-day, day-of-week, and month of application submission. In Panels A and C, a fixed effect for the borrower's sequential loan number is also included. The borrower controls added in column (3) are age, age squared, sex, marital status, number of dependents, log income, and credit score.

Table A6: **Heterogeneity in repayment effects**

Dependent variable: Repayment	(1)	(2)
	A: Marital Status	
	Single/Divorced/Widowed	**Married**
PostBatch	0.012	0.108
	(0.032)	(0.037)
Estimate p-value	0.708	0.003
Bias-corrected estimate p-value	0.833	0.007
Effect as % of pre-cutoff mean	2%	15%
Optimal bandwidth (mins)	[142,133]	[132,101]
Observations within bandwidth	4,054	3,447
Total observations	5,903	5,609
	B: Income	
	Below median	**Above median**
PostBatch	0.028	0.079
	(0.035)	(0.031)
Estimate p-value	0.424	0.011
Bias-corrected estimate p-value	0.492	0.023
Effect as % of pre-cutoff mean	4%	11%
Optimal bandwidth (mins)	[142,121]	[148,121]
Observations within bandwidth	4,017	4,017
Total observations	5,876	5,876
	C: Credit Score	
	None/Marginal/Average	**Better/Best**
PostBatch	0.033	0.142
	(0.028)	(0.044)
Estimate p-value	0.251	0.001
Bias-corrected estimate p-value	0.333	0.004
Effect as % of pre-cutoff mean	5%	18%
Optimal bandwidth (mins)	[148,125]	[117,81]
Observations within bandwidth	5,599	1,834
Total observations	8,140	3,372

Notes: All estimated discontinuities are from linear models that exclude the LBC loan, and loans received within 20 minutes after the LBC. All estimates are from specifications with a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. Heteroskedasticity-robust standard errors of the linear estimates are shown in parentheses below the estimates, calculated using the nearest-neighbor variance estimator with a minimum of three matches. We report both the heteroskedasticity-robust p-values of the linear estimates, and the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates. We also report the estimated effect as a percentage of the pre-cutoff mean repayment rate within the two-hour bandwidth. The optimal bandwidths –rounded to the nearest integer– are reported, observations within the used bandwidth are reported below, and all observations within twelve hours of an LBC below that. All estimates feature fixed effects for the hour-of-day, day-of-week, month of application submission, and the borrower's sequential loan number. All estimates feature controls for age, age squared, sex, marital status, number of dependents, log income, and credit score. These controls drop out when they are the heterogeneous variable of interest.

Table A7: **Heterogeneity by application time**

Dependent variables:	**Application time of day**	
	Before Noon	**After Noon**
	(1)	(2)
A: Induced delay (hrs)		
PostBatch	4.883	12.217
	(1.218)	(1.596)
Pre-cutoff mean	8.685	9.944
Effect as % of pre-cutoff mean	56%	123%
Optimal bandwidth	[165, 94]	[84, 56]
Observations within bandwidth	2,474	3,146
Total observations	3,806	7,706
B: Same-day delivery		
PostBatch	-0.097	-0.292
	(0.032)	(0.036)
Pre-cutoff mean	0.843	0.744
Effect as % of pre-cutoff mean	11%	39%
Optimal bandwidth	[153, 96]	[83, 59]
Observations within bandwidth	2,385	3,173
Total observations	3,806	7,706
C: Repayment		
PostBatch	0.028	0.070
	(0.045)	(0.030)
Pre-cutoff mean	0.728	0.724
Effect as % of pre-cutoff mean	4%	10%
Optimal bandwidth	[128, 110]	[124, 113]
Observations within bandwidth	2,233	4,966
Total observations	3,806	7,706

Notes: Dependent variables: Loan Delay (in hours, panel A); Whether loan was disbursed after the application day (panel B); whether the loan was paid (panel C). Column 1 includes applications submitted between 0.00 hrs and 11.59 hrs. Column 2 includes applications submitted between 12.00 hrs and 23.59 hrs. Day-of-week, hour-of-day, month FEs included, as well as borrower characteristic controls.

Table A8: **Gender and marital status**

Dependent variable: Repayment	Gender of applicant	
	Men	**Women**
	(1)	(2)
A: All		
PostBatch	0.065	0.050
	(0.031)	(0.035)
Pre-cutoff mean	0.715	0.737
Effect as % of pre-cutoff mean	9%	7%
Optimal bandwidth	[159, 120]	[143, 105]
Observations within bandwidth	4,299	3,448
Total observations	6,304	5,208
B: Married sample		
PostBatch	0.072	0.183
	(0.046)	(0.058)
Pre-cutoff mean	0.717	0.706
Effect as % of pre-cutoff mean	10%	26%
Optimal bandwidth	[135, 125]	[132, 77]
Observations within bandwidth	2,263	1,245
Total observations	3,477	2,132
C: Unmarried sample		
PostBatch	0.060	-0.031
	(0.047)	(0.043)
Pre-cutoff mean	0.713	0.760
Effect as % of pre-cutoff mean	8%	4%
Optimal bandwidth	[137, 138]	[145, 123]
Observations within bandwidth	1,906	2,139
Total observations	2,827	3,076

Notes: Dependent variable is whether loan was paid. Day-of-week, hour-of-day, month FEs included, as well as controls for borrower characteristics.

B Data construction

B.1 Batching identification

We do not observe batches' disbursement times (t_2 and t_4 in Figure 1). Instead, we construct the batches and batch cutoffs from our data on disbursement and submission times. Because loans are disbursed from a batch, we observe a series of loan disbursements in quick succession to one another. For example, the median gap between any loan disbursement and the nearest other disbursement in our sample is six seconds, and 94% of loans are disbursed within a minute of another loan. Some loans are processed in isolation of others and appear "unbatched" (not belonging to any particular batch), we exclude these from the data. In particular, our exclusion criteria is to drop all loans that are not disbursed within 2.5 minutes of other loans, as the detectable density of loans falls sharply at that cutoff (see Appendix Figure A6). This drops 259 loans. Alternative approaches, including using k-means clustering to identify batched versus unbatched loans using the minimum distance to another loan produces similar results.

Figure A5: **Distribution of loan application submissions and loan disbursements**

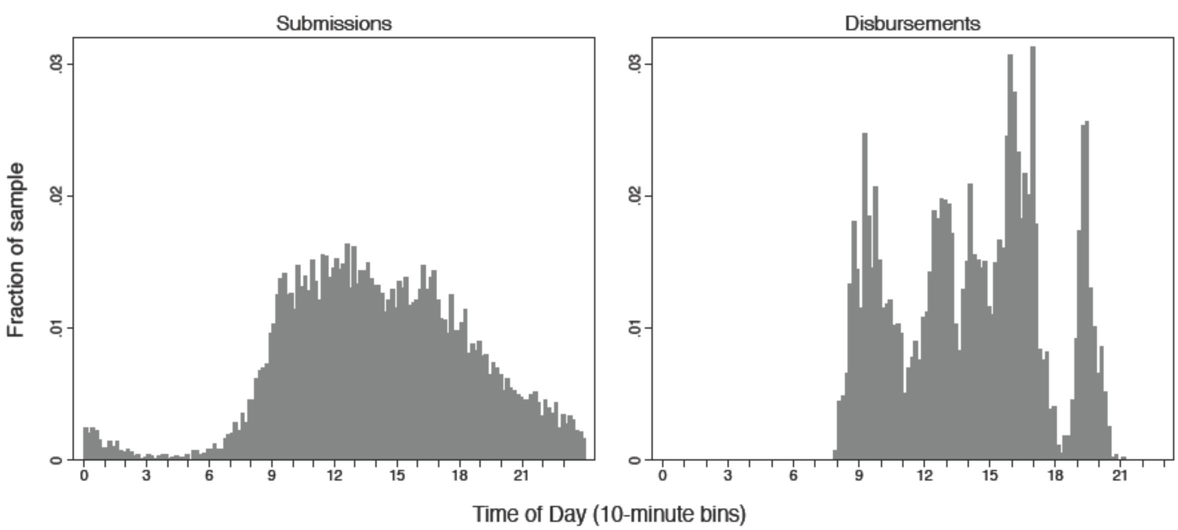

Among the "batched" loans, we use the k-means clustering algorithm to assign each loan to a specific batch within a given day. There are two parts of this process. First, we assume

Figure A6: **Time between loan deliveries in a batch**

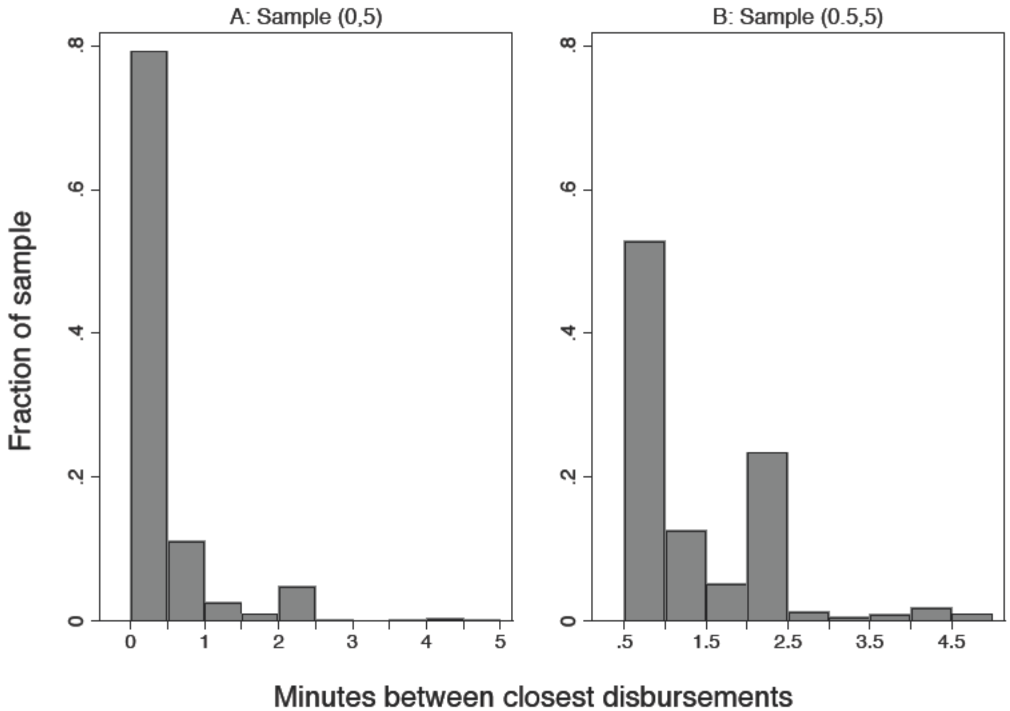

that each day consists of six batches, and let the algorithm assign each loan to one of the six batches under that assumption. Then, we repeat that process assuming five, four, three, two, and one batches per day. Next, we use the maximum proportional reduction of error (PRE) statistic to select the optimal number of batches for each day (Makles, 2012). Of the 142 days in our sample, one is a 1-batch day, 37 are 2-batch days, 47 are 3-batch days, 38 are 4-batch days, eleven are 5-batch days, and eight are 6-batch days.

On some days with smaller samples, random initial batch assignments lead to final batch assignments that overlap, likely representing a locally –rather than globally– optimal assignment. As such, we initialize the algorithm on each day by assigning every k^{th} observation to one of k batches in sequential fashion. This is essentially a stratified randomization procedure to ensure a neutral starting point even in a small sample. In four of 142 days, we still end up with overlapping batches with this approach. By switching to a segmented initial batch assignment, whereby the first N/k observations are assigned to batch one, and the second N/k observations are assigned to batch two, etc. (when assigning N observations to k clusters), we extract non-overlapping batches for these four days (although this initialization does much worse on the overall sample). On one of 142 days, the data clearly suggest one batch is appropriate, but the cluster optimization procedure cannot return an answer of one, so we manually assign all observations on this day to a single batch.

Appendix Figure A7 shows an example of the batching process applied to December 14, 2018. One loan on this day was processed more than 2.5 minutes from the other loans, and is thus removed from the data (left panel). The clustering, applied to the remaining loans, produces four distinct batches around 9am, 12pm, 3pm, and 6pm (right panel).

Figure A7: **Example of batching process**

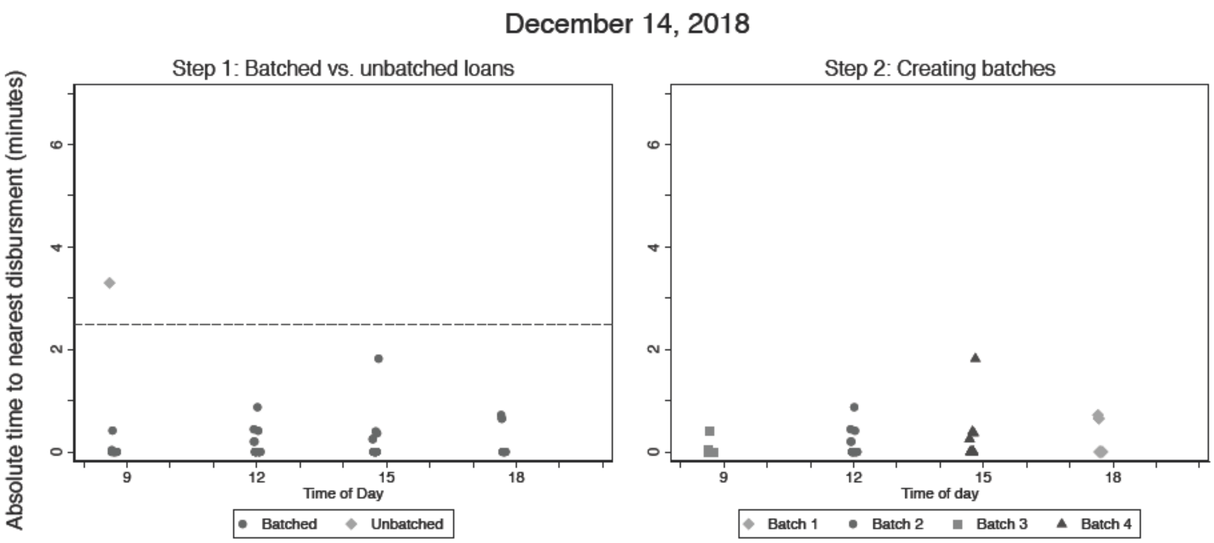

Notes: The left panel shows the first step of our batching process, where we drop one loan that was disbursed 2.5 or more minutes apart from any other loan. Remaining "batched" loans are fed to the k-means clustering algorithm. The right panel shows the batching results from the procedure.

B.2 Constructing the cutoffs

The lower-bound cutoff (LBC) is defined as the latest application submission time within a batch. Appendix Figure A8 shows the distribution of the LBCs in our sample.

We code $DistanceToBatch$ as the difference (in minutes) between loan application submission time and the relevant cutoff, and $PostBatch$ as an indicator for whether $DistanceToBatch$

Figure A8: **Distribution of lower-bound cutoffs**

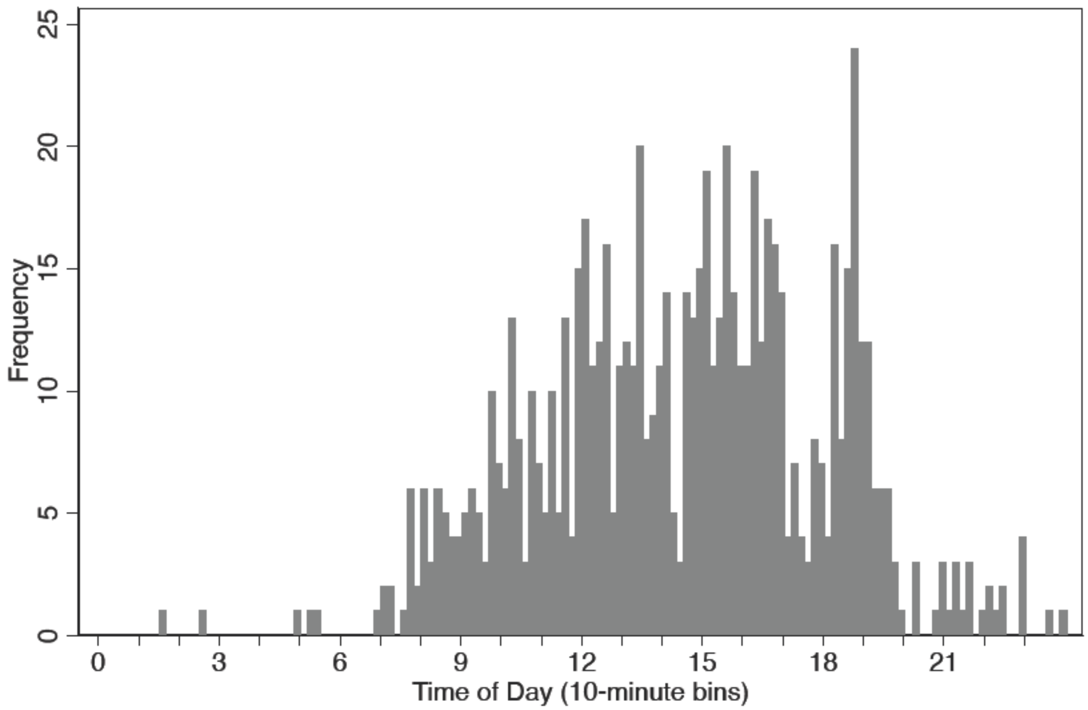

is positive. Accordingly, in Figure 1, the application submission (start of the verification process) for loan k is closer to the Batch A cutoff than to the Batch B cutoff. Therefore, loan k is assigned to the Batch A cutoff, with a positive value of $DistanceToBatch$ ($PostBatch = 1$). The start of the verification for loan m is closer to the Batch B cutoff. Therefore, loan m is assigned to the Batch B cutoff, with a negative value of $DistanceToBatch$ ($PostBatch = 0$).

B.3 Selection around the cutoffs

Because the density of submissions drops at the LBC, the typical regression discontinuity validity test–smoothness of the density of the running variable–is not informative in our context. Appendix Figure A3 shows the density of loan submissions within a 4-hour window of an LBC, excluding the LBC loan. Note that, because every window around an LBC must, by definition, contain a submission for which $DistanceToBatch = 0$, zero is overrepresented in the distribution. We exclude those observations from the figure. It shows a very clear violation of smoothness at the LBC.

Additionally, there is a selection issue for both LBC loans and loans submitted just after the LBC. It is possible that loans after the LBC were processed in subsequent batches because they were more difficult to process; if they had been easy to process, they would have been included in the batch with the LBC loan, and become the LBCs themselves. If processing difficulty is negatively correlated with borrower quality, a failure to fix these issues could lead to biased estimates of the β_2 coefficient in Equation (1) towards indicating harmful effects of induced delays. For example, a failure of the borrower to pick up the phone the first time they are called for identity verification could be correlated with borrower quality. The average time from submission to disbursement is 19.6 hours for loans submitted within 20 minutes after the LBC, and 17.7 hours for loans submitted 20-60 minutes after the LBC. This supports the idea that loans right after the LBC take longer to process, and that they could be negatively selected.

To determine how to exclude these loans, we first consider the smoothness of the density of the running variable above the LBC. We use the "rddensity" suite of commands developed by Cattaneo et al. (2018) to determine where the right side of the density shown in Appendix Figure A3 achieves smoothness, starting from the LBC; where does it shift from outlier loans that couldn't be processed quickly enough to be the LBC to typical loans that simply missed the previous batch? Starting at five minutes post-LBC, we test for smoothness through each five-minute increment above the LBC, up to one hour. We use the optimal bandwidth approach, with bias-correction robust standard errors. Appendix Table A9 shows the p-value associated with each test, along with the optimal bandwidth and effective observation count.

The first failure to reject is at 15-minutes post LBC, although the estimated optimal bandwidth exceeds the given range (a test with a symmetric bandwidth of just under 15 minutes yields a p-value of 0.175). Beginning with 20-minutes post-LBC, we always reject the null with a well-defined bandwidth.

Does a 20-minute exclusion window make sense using other approaches? We now test

Table A9: **Density-smoothness tests of post-LBC application submissions**

Minutes post-LBC	p-value	Optimal bandwidth	Obs. in bandwidth
5	0.003	[3,41]	1,276
10	0.050	[6,44]	1,508
15	0.228	[25*,49]	1,971
20	0.380	[11,60]	2,281
25	0.805	[9,78]	2,723
30	0.922	[18,132]	3,892
35	0.257	[12,103]	3,213
40	0.144	[13,77]	2,706
45	0.447	[23,73]	2,911
50	0.279	[20,70]	2,673
55	0.745	[19,67]	2,468
60	0.577	[26,64]	2,595

Notes: * this bandwidth is outside the range of the data. Optimal bandwidths are rounded to the nearest integer. Discontinuities are estimated with a quadratic fit of the density and a triangular kernel. We use distinct optimal bandwidths left and right of the cutoffs to allow for the larger amount of data to the right of these cutoffs to improve precision. p-values are from the heteroskedasticity and bias-correction robust standard errors, calculated using the nearest-neighbor variance estimator with a minimum of three matches.

directly for smoothness in a key observable –creditworthiness– through post-LBC cutoffs. Using the rdrobust optimal bandwidth approach, in Table A10, we show how the assessed credit score category of borrowers changes when a batch cutoff is missed. Of the 65 p-values in the table, only two are less than 0.05, and both of these are associated with fewer "best" score borrowers being in the sample after the LBC –consistent with our concern regarding negative selection in this period right after the LBC. Focusing on that credit score category, the discontinuities at zero, five, ten, and 15 minutes post-LBC are at least marginally statistically significant, and we fail to reject smoothness at twenty minutes.

Table A10: **Borrower credit score smoothness through post-LBC cutoffs**

Credit score:	None		Marginal		Average		Better		Best	
Minutes post-LBC	Coef.	p-value	Coef.	p-value	Coef.	p-value	Coef.	p-value	Coef.	p-value
0	-0.004	0.834	-0.033	0.187	0.038	0.141	0.011	0.559	-0.028	0.015
5	-0.009	0.627	-0.033	0.229	0.041	0.133	0.008	0.708	-0.024	0.061
10	-0.005	0.809	-0.025	0.364	0.025	0.366	0.013	0.543	-0.027	0.035
15	0.006	0.713	-0.024	0.401	0.025	0.393	0.005	0.843	-0.023	0.090
20	-0.006	0.841	-0.003	0.935	0.011	0.807	0.009	0.728	-0.017	0.256
25	-0.006	0.837	-0.003	0.995	0.002	0.898	0.014	0.552	-0.015	0.310
30	-0.003	0.913	-0.005	0.976	0.004	0.941	0.010	0.717	-0.012	0.441
35	-0.011	0.501	0.010	0.490	0.005	0.920	0.013	0.661	-0.022	0.091
40	-0.018	0.219	0.019	0.288	0.011	0.705	0.005	0.894	-0.021	0.080
45	-0.019	0.153	0.005	0.713	0.012	0.603	0.016	0.456	-0.020	0.105
50	-0.022	0.071	-0.001	0.995	0.011	0.567	0.024	0.215	-0.020	0.143
55	-0.020	0.078	0.003	0.998	0.009	0.445	0.015	0.480	-0.019	0.255
60	-0.012	0.119	0.001	0.631	0.006	0.428	0.018	0.350	-0.019	0.424

Notes: Estimates exclude loans received between the LBC loan and the minutes post-LBC. All specifications use a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. p-values are from the heteroskedasticity-= and bias-correction robust standard errors, calculated using the nearest-neighbor variance estimator with a minimum of three matches.

Finally, in Table A11, we test for whether observable borrower characteristics are smooth through the 20-minute post-LBC cutoff, using the rdrobust optimal bandwidth approach. Note that while these borrower characteristics are fixed at the individual level (we only observe loan amount and length for the first loan), the unit of observation is the loan: a particular borrower can experience both sides of the cutoff. Therefore, we use the full sample of loans. Since we fail to measure any significant or large jump at the cutoff for any variables, we use the 20-minute post-LBC latent cutoff as our preferred specification.

Table A11: **Borrower variable smoothness at 20-minute post-LBC cutoff**

N = 11,512	Coef.	S.E.	p-value	Effect size	Optimal BW	Obs. in BW
Age	0.525	0.534	0.341	1%	[148,98]	7,436
Female	-0.008	0.030	0.989	-2%	[119,102]	6,828
Married	-0.002	0.035	0.787	-0%	[92,144]	6,641
Dependents	-0.053	0.081	0.416	-4%	[90,109]	6,077
Log income	0.003	0.040	0.952	0%	[142,108]	7,485
Credit score	0.010	0.065	0.966	1%	[115,129]	7,182
Loan amount	-7.799	20.847	0.589	-0%	[124,122]	7,275
Loan length	0.217	0.415	0.556	1%	[123,133]	7,447

Notes: Estimates exclude the LBC loan, and loans received within 20 minutes after the LBC. All specifications use a triangular estimation kernel, and an optimal bandwidth selected from a 12-hour window around the LBC. Heteroskedasticity-robust standard errors (calculated using the nearest-neighbor variance estimator with a minimum of three matches) are shown. We also report the bias-correction- and heteroskedasticity-robust p-values of the quadratic, bias-corrected estimates. The reports effect size is as a percentage of the pre-cutoff mean value of the borrower characteristic within the two-hour bandwidth. The optimal bandwidths (rounded to the nearest integer) are reported along with observations within the used bandwidth. The overall sample size for all models in the table corresponds to all loans within twelve hours of an LBC.

One test that we do not report in full is the test for density smoothness of the running variable $DistanceToBatch$ with loans right after the LBC excluded. This is because failure to reject here can come either from densities that match up nicely at the post-LBC cutoff, or from an increase in the standard error in the exclusion window to the left of the cutoff. Without data in the region around the cutoff, the uncertainty about the density is large. In the main analysis, this simply reduces our power. However, in this analysis, where we are seeking a region where we cannot reject smoothness, it may lead us to be too confident in the selection of a smaller exclusion window.

C Impact of delays on lender

In this appendix we explore the effects of delays on the welfare of the firms. For this discussion, we first note that we do not have information on the cost side of the firm, limiting our ability to quantify the impact on firm profits.

First, note that the primary cost associated with a loan that has been issued is following up with collection in case the loan is delinquent. Assuming that the lender pursues delinquent loans only if doing so is profitable, it is likely that the increased repayment of delayed loans increases the firms' profits.

Next, in addition to the positive, direct effect of the delay on the profits of the firm, there are indirect effects to consider–namely, the knock-on effects of delays on the clients' demand for future loans. There are two potentially countering effects of the delay on demand. One the one hand, the lender offers additional loans only to clients who have not in arrears. Because delayed loans are repaid more often, the lender increases its potential client base. On the other hand, borrowers whose loans have been delayed might be less willing to borrow again from this lender, as the delay could be construed as a signal of low lender quality (or simply a lender that is too "slow".) This could depress the demand. Finally, note that borrowers that are induced to repay and then borrow again might be "marginal" borrowers, i.e., they have a high propensity to default on future loans.[21]

We study these effects by measuring the effect of being delayed on the first loan on the likelihood of borrowing again. In our analysis, we limit the sample to first-time borrowers who fall within the 2-hour batch disbursal window, and estimate equation (1) on the following outcome variables: whether the client borrowed at least one more time; whether the borrower repaid the second loan (conditional on borrowing a second time); and the total number of loans the borrower obtained from the lender. In these regressions, $PostBatch_i$ is an indicator for the borrower having missed the cutoff time at the time of the *first* loan application,

[21]We do not consider here other second-order effects of delays, such as the possibility that they negatively impact the reputation of the lender and make acquisition of new clients more expensive.

and the set of controls X_i include $DistanceToBatch$ (the number of minutes between the submission time of the first loan and the LBC loan); day-of-week, hour of day and month fixed effects; and borrower controls.

Estimates of β are presented in the table below, with and without controls. None of the estimates are statistically significant, although all the estimates are positive. Delays are not affecting the likelihood that borrowers take on credit in the future (columns 1 and 2), that they repay their second loan if they borrow again (columns 3 and 4), or the total number of loans (columns 5 and 6). It is possible that the lack of statistically significant results is due to the limited number of observations; note that the p-value for the total number of loans is 0.11, which is close to significant at conventional levels.

Table D1: **Effect of delays on first loans on subsequent loan outcomes**

Dependent Variables:	Borrowed again		Repaid second loan		Total number of loans	
	(1)	(2)	(3)	(4)	(5)	(6)
$PostBatch$ (first loan)	0.038	0.048	0.035	0.051	0.088	0.105
	(0.038)	(0.038)	(0.048)	(0.048)	(0.069)	(0.067)
Obs. within bandwidth	3,478	3,424	1,898	1,870	3,483	3,487
Estimate p-value	0.310	0.201	0.459	0.287	0.198	0.119
Controls	No	Yes	No	Yes	No	Yes
Day-of-week f.e.	No	Yes	No	Yes	No	Yes
Hour-of-day f.e.	No	Yes	No	Yes	No	Yes
Month f.e.	No	Yes	No	Yes	No	Yes

Notes: RD regressions at the borrower level, following equation (1). Controls include: age, sex, income, marital status, number of dependents, credit score. See Table 1 for more information on estimation. Heteroskedasticity-robust standard errors in parentheses.