

Too Fast, Too Furious?

Digital Credit Delivery Speed and Repayment Rates[§]

Alfredo Burlando[¶] Michael A. Kuhn^{||} and Silvia Prina^{**}

February 17, 2023

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 findings suggest that selectively slowing down credit could improve lender profitability and help consumers avoid default.

JEL Classifications: D14, D18, G51, O16

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

[§]We are grateful to Carlos Chiapa, Bilge Erten, Benjamin Hansen, Glen Waddell, Dilip Mookherjee, Tavneet Suri, Bob Triest, Judy Chevalier, Kevin Lang, Diego Ubfal, Prashant Bharadwaj, David Freeman Niels Hermes and seminar participants at ASSA, Bocconi University, CEGA, DCO, NEUDC, Simon Fraser University and the University of Padua for helpful comments. We thank our digital lender partner for providing us with the data and answering our many questions. The authors thank the Digital Credit Observatory for generous research support. Any errors are our own. The dataset analyzed in the current study is not publicly available because it is proprietary to our lender partner, but is available from the corresponding author on reasonable request. All code is available on request.

[¶]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). For example, by 2018 over a quarter of Kenyan adults had an active digital loan (Totolo, 2018) and (Björkegren et al., 2022) report the existence of over 50 digital lenders in Nigeria in contrast to limited availability of formal financial services. Data harvesting and analysis have enabled digital credit providers to assess consumers' creditworthiness and ability to repay without requiring any collateral to secure loans (Björkegren and Grissen, 2018). 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, Björkegren et al. (2022) find that it increases subjective well-being, and Brailovskaya, Dupas, and Robinson (2021) find that increases perceived financial well-being.

On the other hand, many borrowers struggle to repay digital loans (Carlson, 2017).¹ 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).² 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); Brailovskaya et al. (2021)) and may use them to finance unproductive, time-sensitive investment and consumption opportunities such as gambling (Malingha, 2019). Given that the industry suffers from high default rates, policy makers have started to advocate for consumer protection measures targeting the digital credit market (Donovan and Park, 2019).

¹The default rate is 27% in our data from Mexico. Other studies have found default rates of 15% (Brailovskaya et al. (2021) in Malawi) and 7% (Bharadwaj et al. (2019) in Kenya and Björkegren et al. (2022) in Nigeria).

²Evidence from the credit card market shows that less-sophisticated borrowers may be susceptible to over-borrowing, penalties, and backloading repayments (Meier and Sprenger, 2010; Heidhues and Kőszegi, 2010).

To date, the speed of digital credit delivery has not been studied.³ We address this knowledge gap with a unique administrative dataset of digital loans and quasi-experimental variation in the time that elapses between the approval of a loan request and when the loan is deposited in the borrower’s bank account. Our data consist of loan records from the full set of approved clients over a seven-month period in 2018-2019 from a digital lender operating in Mexico. These records include both loan approval and disbursement timestamps, which we use to measure loan delivery speeds.

We cannot naively relate our observed delivery speeds to repayments due to the fact that processing times may be correlated with the *ex-ante* propensity to repay a loan. For example, borrowers who make mistakes in their loan application might face longer loan delivery times if that mistake needs to be fixed prior to disbursement. Such borrowers might have characteristics that correlate with low repayment rates. We thus rely on a quasi-experimental approach that takes advantage of variation in delivery speeds coming from the fact that the lender disburses loans in batches about two to four times per day. Loans added first to a new batch remain in the queue longer than those added last, leading to systematic differences in processing times between loans. Our empirical strategy uses 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 precisely time their applications for faster service. However, unlike the standard regression discontinuity (RD) setup, we do not observe the precise moment when an application cannot be processed fast enough to enter the current batch; 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

³We are careful to distinguish delivery speed from the speed of digital credit demand –the period in between multiple loans from the same borrower– which is studied by Carlson (2017).

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. This point estimate translates to a 21% reduction in the likelihood of loan default when loan delivery is slowed down. Our results are in stark contrast with the OLS regression estimates showing a significant *negative* correlation between loan processing times and repayment rates. While the welfare effects of credit delays for borrowers are ambiguous and depend on the mechanism through which they operate on repayment, we find that delays increase the lender’s per-loan profit without decreasing loan demand (at least in the short run). These findings suggest that in a high-default environment, waiting periods can be a valuable addition to the digital lender’s toolkit.

How can a disbursement friction *increase* loan repayment? While our data do not allow us to conclusively identify a single mechanism, we evaluate a variety of plausible possibilities and suggest a mechanism that best fits the data. Using the terminology of Karlan and Zinman (2009), we group potential mechanisms under the *hidden information* and *hidden action* umbrellas. In the case of hidden information, an extra delay causes selection out of borrowing that improves the lender’s risk within the remaining pool. We rule out this potential mechanism for two reasons. First, in our sample borrowers cannot easily select out of a loan once it is sent for processing. Hence, such cases are extremely rare. Second, there is no effect of an extra delay on the rate of selecting out.

Considering hidden action, an extra delay causes borrowers to modify their behavior in a way that increases their probability of repayment. Under the hidden action umbrella, we outline four separate channels. The one that best fits our data best is what we refer to as the *behavior channel*: following an unexpected delay, borrowers switch their preferred use of loans in a way that improves the chances of repayment. This happens either because the delay changes borrower’s preferences (e.g. a borrower wishes to use the funds to gamble or to purchase temptation goods but the delay causes her to reconsider this plan), or because the delay changes who makes the decision about how to use the loan (e.g. a borrower takes

out a loan on her own, but the delay leads her to discuss the loan with her partner before the funds arrive). Household bargaining may reduce the likelihood that the borrower makes negative-return investments (e.g. gambling) or spends the funds on temptation goods. In this case, repayment could improve even if the funds are simply left unused, but in general, if the delay increases the expected return on loan funds, repayment will increase due to changes in borrower behavior. We find evidence supporting this hypothesis: the effects on repayment are driven by married borrowers whose loans are delayed to the next day.

Our results on the importance of delays 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 (or “cooling-off”) periods are typically considered in the context of impulsive decision-making, where the choice set is not materially affected by the waiting period. They are 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 is also connected to the literature on behavioral biases in consumer financial choice, where restrictions on choice sets have the potential to be welfare-improving. Behavioral biases induce people to engage in sub-optimal 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 people’s choices.⁴ There is a subset of this literature on the role of behavioral biases in high-interest lending, which include digital loans but also payday loans in the U.S. Alcott, Kim, Taubinsky, and Zinman (2021) find, in the context of payday loans in the U.S., that borrowers are willing

⁴Examples include commitment savings accounts (Ashraf, Karlan, and Yin, 2006) and frequent fixed payments for microfinance borrowers (Bauer, Chytilová, and Morduch, 2012; Field, Pande, Papp, and Rigol, 2013).

to pay a significant premium to avoid future borrowing. Bertrand and Morse (2011) show that psychology-guided information disclosure reduces the use of payday loans.

Our results are also related to a microfinance literature studying the effects of product characteristics on repayment behavior. Karlan and Zinman (2009) find a 21% reduction in loans in collection status (2.5 percentage points) when borrowers are offered dynamic incentives. Similarly, Karlan et al. (2015) find a 27% reduction in unpaid loans 30 days past-due (3.7 percentage points) when borrowers are sent text message reminders; this study’s light-touch intervention makes it the best comparison available for the effect of disbursement delay. Field et al. (2013) find that providing a repayment grace period reduced repayments by 6 percentage points, which is a 370% increase in the default rate, while Feigenberg et al. (2013) finds that more frequent microfinance group meetings increase repayments by 5.1 percentage points, which corresponds to a 72% decrease in the default rate. Thus, our finding of a 21% reduction (5.6 percentage points) in default is not an outlier in this literature.

The paper proceeds as follows. Section 2 describes the setting, sample, and key variables. Section 3 explains the empirical strategy. Section 4 presents the results and Section 5 provides a discussion of potential mechanisms. Section 6 discusses policy implications and concludes.

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),⁵ and the loan terms vary from seven to 30 days. The cost of a loan (and thus repayment amount) is fixed at the time it is taken; early payment is allowed at any time, but does not save the borrower any interest. Costs consist of interest, taxes, and fees, with the implied APRs reaching up to 478.8%. There is a one-time penalty added to the amount owed for delinquent loans. The characteristics of this loan product are comparable to those of other digital lenders in the

⁵The exchange rate during the study period is approximately USD 1= MXP 20.

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. For repeat borrowers the lender also considers their prior repayment behavior, rejecting applicants who failed to fully pay back previous loans.

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. For first-time borrowers, processing includes an identity verification 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.

2.2 Sample

Our sample consists of 11,512 approved loan applications from 7,206 borrowers, with loans disbursed between November 2018 and May 2019.⁶ Forty-eight percent of the loans in our

⁶The 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 3.2 and 3.3 detail the additional steps to determine the estimation sample.

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 ‘delays’ are calculated as the time (in hours) between loan application submission by the client and disbursement by the bank. First-time borrowers face a delay of 24 hours on average (median 18 hours), while for repeat borrowers it is 9 hours (median 3 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%.⁷ 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, 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.

⁷Unfortunately, we cannot tell whether a loan has been partially repaid. While we observe repayments through November 11, 2019 (164 days after the last loan), it is possible that some of the defaulted loans were repaid after we received the data. However, of the loans issued through February 2019 which were repaid, 98% were paid before 164 days.

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, on any given day, there are no set times at which batches are sent to the bank for disbursement.⁸ However, first we have identify the batches and batch cutoff times within our data. Here we outline our procedures and refer to Appendix Section B for additional details.

3.1 Constructing batch cutoffs for regression discontinuity


Figure 1 shows a simplified timeline of loan applications and disbursements to illustrate our approach to identification. We will refer to this figure throughout this section and define each of its relevant features in turn. Borrowers in our sample are approved for their loans at the end of their online application, at which point their loans are submitted for verification processing. This period is represented by the rectangles associated with each loan in Figure 1, where the left edge of the rectangle represents the loan submission time, and the right edge represents the processing completion time.

Once processed, loans are assigned to the open disbursement batch. For example, loans a and b are both assigned to Batch A because processing is finished prior to disbursement time t_2 . At t_2 , Batch A is disbursed and the queue is cleared. These Batch A loans are marked with blue outlines in Figure 1. Loans c , d , and e are not processed until after t_2 , but before the Batch B disbursement at t_4 . They are thus assigned to the next disbursement batch, Batch B. Batch B loans are marked with red outlines.

To make it into batch A, a loan must be processed before disbursement time t_2 . We thus define Batch A Cutoff time t_1 as the last moment in time when a borrower could have

⁸The lender is not aware of these batch cutoff times in advance either, and they exhibit significant variability. Appendix Figure A1 shows the distributions of submissions and disbursements throughout the day.

Figure 1: Hypothetical timeline of loan submission, processing 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 3.2.

submitted a loan request and have some probability that it could be processed prior to t_2 . Since the Batch A cutoff is unobserved, we proxy for it using the last submission time among all loans in Batch A. In Figure 1, this is loan *b*. We refer to this proxy as the Batch A lower-bound cutoff (LBC), and in Figure 1 we distinguish between the unobserved Batch A true cutoff at t_1 , and its proxy the Batch A LBC.

Our empirical strategy is best illustrated by the comparison between loans *a* and *d*. These loans have been submitted by two separate clients around the same time and take a similar amount of time to be processed. However, their submission times fall on different sides of the Batch A LBC, and thus loan *d* has a longer delay than loan *a*.

To implement this strategy, we first assign every loan to the closest batch LBC (based

on its application submission time).⁹ This is the Batch A LBC for all loans in Figure 1. Next, we create an indicator called *PostBatch* that takes a value of one if the application was submitted after its assigned LBC. In Figure 1 loans assigned to the pre-cutoff period are shaded in orange (loans *a*, *b*, and *e*) and loans assigned to the post-cutoff period are shaded in green (loans *c* and *d*; we explain why loan *c* is shaded as a mix of green and orange in the next section). Finally, we compute a continuous variable labeled *DistanceToBatch* that represents the time of loan application submission minus the assigned batch cutoff time. Loans for repeat borrowers follow a different processing procedure from loans for first-time borrowers, and as a result the minimum potential processing time for repeat loans is much shorter.¹⁰ Thus, we calculate separate LBCs for first-time and repeat borrowers.

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

$$\begin{aligned}
 Y_{ij} = & \beta_1 \text{DistanceToBatch}_{ij} + \beta_2 \text{PostBatch}_{ij} + \\
 & \beta_3 \text{DistanceToBatch}_{ij} \times \text{PostBatch}_{ij} + \delta X_{ij} + \epsilon_{ij}
 \end{aligned} \tag{1}$$

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

⁹Recall that on 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.


¹⁰LBC loans are processed in 113 minutes on average for repeat loans whereas first-time LBC loans are processed in 490 minutes.

figures. Because *PostBatch* is assigned at the loan level, we do not cluster standard errors.¹¹

3.2 Constructing batches

We do not explicitly observe the batch to which a loan is assigned. In Figure 1, this means that we do not observe the batch disbursement times t_2 and t_4 . However, 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. For this step, we pool the records of first-time and repeat borrowers because they are entered into shared disbursement queues. See Appendix Section B.1 for additional details on this process.¹²

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.

¹¹See Abadie, Athey, Imbens, and Wooldridge (2017).

¹²The k-means algorithm we use to construct batches is initialized in a deterministic fashion. While inspection of the resulting batches suggests the deterministic procedure is superior, in Appendix Section B.1 we report the results of an approach where we instead use a stochastic initialization and simulate the batching process 1000 times to produce a distribution of our main estimations. Results are in Appendix Figure B3 and show that in most cases our main estimates in the paper are conservative relative to the distribution.

Figure 2 shows the results of our procedure by plotting the likelihood a loan is 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 A2 and A3). 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.

3.3 Challenges to identification

Because we do not observe the true batch cutoffs there is a potential for bias in our empirical approach if the amount of time it takes to process a loan depends in part on factors related to borrower quality. In this section we carefully review the biases that stem from this issue, and outline the solutions we use to correct for them. We also provide simulation evidence that whatever bias we cannot eliminate will only lead us to underestimate the effect of delays, but that the magnitude of this bias is likely very small.

Loans *b*, *c*, and *e* in Figure 1 are drawn to help us illustrate our key challenges to identification and exactly why these problems stem from the unobservability of true cutoffs. Note that submission prior to the Batch A cutoff is not sufficient for a loan to be in Batch A; consider loans *c* and *e* in Figure 1. Both are submitted prior to the Batch A cutoff, but because of the time it takes to verify them, they don't make it into Batch A. This fact, combined with the unobservability of the true batch cutoffs, creates the potential for selection bias. This is illustrated by loan *c* in Figure 1. Loan *c* is submitted after the Batch A LBC –and is thus assigned to the Batch A post-cutoff period– but before the true, unobserved

Batch A cutoff, t_1 . If we could observe that cutoff, we would have instead assigned loan c to the Batch A pre-cutoff period. Hence, we shade loan c in both green and orange to represent the fact that it is wrongly assigned to the post-cutoff period. In general, loans submitted right after LBCs that take a long time to process are likely to be wrongly assigned in this way, leading to too few long processing time loans assigned to the pre-cutoff period.

This selection problem can be seen in our data. 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 Table B3 provides additional evidence that loan applications submitted within 20 minutes after the LBC are negatively selected along observables.

This same process leads to a second selection problem: imagine if loan c in Figure 1 had been processed very quickly and added to the queue prior to the Batch A disbursement at t_2 . In this case, loan c would supplant loan b as the LBC loan. Thus, while loans that appear in our data right after LBCs are likely to feature long processing times, the LBC loans themselves –e.g. loan b in Figure 1– are likely to feature short processing times. Indeed, 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.

In addition to these selection problems, using LBCs as batch cutoff proxies creates two discontinuities in the density of application submissions –the smoothness of which is a key test for the validity of RD designs. First, LBC loans are vastly over-represented in the data. By construction, every single batch cutoff window in our data features a loan with $DistanceToBatch = 0$: the LBC loan. On the other hand, the density of loan submissions prior to the LBC loan represents the natural rate of submission arrival, creating a discontinuous spike in the density of submissions at $DistanceToBatch = 0$ (see Appendix Figure B5, Panel A).

Second, the density of submission times right after an LBC is lower than the density

right before it (see Appendix Figure B5). This is not due to active manipulation by the applicants or the lender; rather it arises mechanically through the process of defining each LBC as the *latest* of all submission times within a batch. A random lull in submissions increases the likelihood that the last submission prior to that lull is the latest submission time within a batch. If the lull hadn't occurred, loans after the LBC would have made it into the same batch and supplanted the LBC (unless such loans were slow to process, generating the selection problem discussed earlier).¹³

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-donut” RD, where we drop only the right side of the donut hole.¹⁴ We arrive at an exclusion window of 20 minutes for our main estimates after a variety of exercises, detailed in Appendix Section B.3. These include: 1) testing for when, after the LBC, the submission density becomes smooth, 2) testing for when, after the LBC, borrower credit score becomes smooth with respect to loans received just prior to the LBC, 3) testing for whether other borrower characteristics are also smooth at the 20-minute post-LBC mark (again with respect to loans received just prior to the LBC), and 4) simulating the amount of selection bias we should expect based on the size of the exclusion window. The first two exercises suggest an exclusion window of 20 minutes, the third confirms that all available borrower observables are smooth within such a window, and the fourth suggests that for almost any exclusion window, the correlation between processing time and repayment likelihood would have to be implausibly large for meaningful selection bias to exist. Indeed, in Section 4.3 we show that our main estimates do not depend heavily on the exact size of the exclusion

¹³This phenomenon is related to what Miller and Sanjurjo (2018) call “streak selection bias” in the context of collecting data to analyze the hot hand fallacy. In both cases, the way in which the researcher constructs an estimation sample within their data induces a selection process in a dimension that should be independent of the true data-generating process.

¹⁴Beyond selection concerns, the half-donut sample also has the added benefit of reducing measurement error associated with proxying for the batch cutoffs with LBCs as long as the exclusion window doesn't overshoot the true cutoff by more than its distance to the LBC.

window (and to the degree that they vary, they do so in exactly the manner predicted by our simulation exercise). Finally, we note that the selection bias that would exist from immediate post-LBC loans is entirely in the opposite direction of our main result that delays increase repayment. These exclusions yield our main estimation sample of 11,512 loans.

Importantly, selection concerns are absent for the loans that were submitted before the LBC because our sample construction process makes no restriction on whether those loans can be processed in either the same batch as the LBC or in following batches, i.e., they are not selected based on their processing speed, nor is there any chance they are erroneously assigned to the pre-cutoff period.¹⁵ We thus include all loans leading up to the LBC in the pre-cutoff period.


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 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.

¹⁵We provide clear evidence that selection problems arise discontinuously at the LBC, rather than as the LBC is approached from the left using both our observed data and a simulation. See Appendix Section B.3

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 ($p < 0.001$ for all models and samples). 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 decrease the likelihood of same-day disbursement. Appendix Table A4 shows that missing a batch cutoff reduces the likelihood that a borrower receives her loan on the same day as her application by 24 percentage points ($p < 0.001$ for all models and samples).

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 while 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;¹⁶ the effect magnitude as a percentage of the pre-cutoff mean repayment within two hours of the cutoff; the optimal bandwidth as determined by the `rdrobust` command; and the number of observations within that optimal bandwidth.¹⁷

¹⁶The 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 selection, but it pertains to the quadratic estimate used for bias correction, not the linear estimate of interest.

¹⁷The 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.

Table 1: Impact of the cutoff on loan repayment

RD bandwidth:	Two-hour		Optimal	
	(1)	(2)	(3)	(4)
A. Full sample (N = 11,512)				
<i>PostBatch</i>	0.059 (0.023)	0.063 (0.024)	0.060 (0.024)	0.056 (0.024)
Estimate <i>p</i> -value	0.011	0.008	0.012	0.018
Bias-corrected estimate <i>p</i> -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)				
<i>PostBatch</i>	0.041 (0.036)	0.040 (0.035)	0.041 (0.034)	0.054 (0.034)
Estimate <i>p</i> -value	0.259	0.251	0.227	0.110
Bias-corrected estimate <i>p</i> -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)				
<i>PostBatch</i>	0.064 (0.030)	0.083 (0.034)	0.078 (0.034)	0.074 (0.034)
Estimate <i>p</i> -value	0.037	0.015	0.021	0.029
Bias-corrected estimate <i>p</i> -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).

For the full sample, our preferred specification in Column (4) shows that the induced delay (10 hours on average) increases repayment rates by a statistically significant 5.6 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. For the sub-samples in Panels B and C, we find a statistically significant 7.4 percentage point (10%) increase in repayment for repeat loans ($p = 0.029$) and a 5.4 percentage point (8%) increase in repayment for first-time loans ($p = 0.110$).

4.3 Robustness tests

Given that we construct both the batches and the cutoffs, we consider the robustness of our estimates in both of these dimensions. First, we verify the robustness of the estimates with respect to the imputation of the batches, described in Appendix Section B.1. Our approach might introduce errors if the K-means clustering approach mis-classifies loans by assigning them to incorrect batches. To evaluate how sensitive our results are to this problem, we generate alternative loan batches by initializing the K-means algorithm with a random partition of the data on each day. We create 1,000 samples this way, allowing us to plot the density of estimates for each of the 12 estimates in Table 1. These densities are in Appendix Figure B3. The 5th percentile of the estimate distribution is positive in 11 out of 12 cases.¹⁸ While our main estimates in Table 1 are not always centered in the distributions, our preferred estimates in column (4) are generally conservative according to this test.¹⁹


Second, we show the robustness of the main estimate with respect to the size of the post-LBC exclusion window (i.e. the size of the RD donut hole). Figure 4 shows that our preferred full-sample estimate from Table 1, column (4), is very stable when we alter the

¹⁸The exception is Panel B, column (2), which does not correspond to a significantly positive in Table 1.

¹⁹The median (and mean) estimate corresponding to Panel A, column (4) from this approach is a 6.7 percentage point increase in repayment, rather than the 5.6 percentage points in Table 1.

window between zero and 40 minutes. The estimate increases slightly and is more precise for larger exclusion windows. This positive relationship between the size of the exclusion window and the size of our estimate is consistent with immediate post-LBC loans being negatively selected on borrower quality. Appendix Figure B6 shows simulations of the relationship between the size of the exclusion window and the amount of selection bias we should expect over a wide range of assumptions about how much unobservable borrower quality varies with loan processing time. We find that the selection bias will always lead us to underestimate the (positive) effect of delays on repayment. However, the correlation between processing time and repayment likelihood would have to be implausibly large for meaningful selection bias to exist.

Figure 4: **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).

4.4 Comparison with OLS estimates

In Appendix Table A6 we report the results of a naive OLS analysis on the full loan dataset, in which we regress loan delays on the repayment indicator. We report the results without controls in column (1), with time controls in column (2), and with borrower characteristics in column (3). In contrast to our RD results, the coefficient estimates on loan delays for the full sample (Panel A) are negative and significant. Column (1) shows that each hour of delay is associated with a 0.1 percentage point decrease in repayment. Controlling for time and borrower characteristics decreases the magnitude of the association to 0.04 percentage points. When considering first-time and repeat borrowers separately (Panels B and C, respectively), the estimates are negative as well, but become statistically insignificant with controls. Overall, the OLS estimates suggest that the endogeneity of delay length may be important, highlighting the necessity of our RD method for delivering a causal estimate.

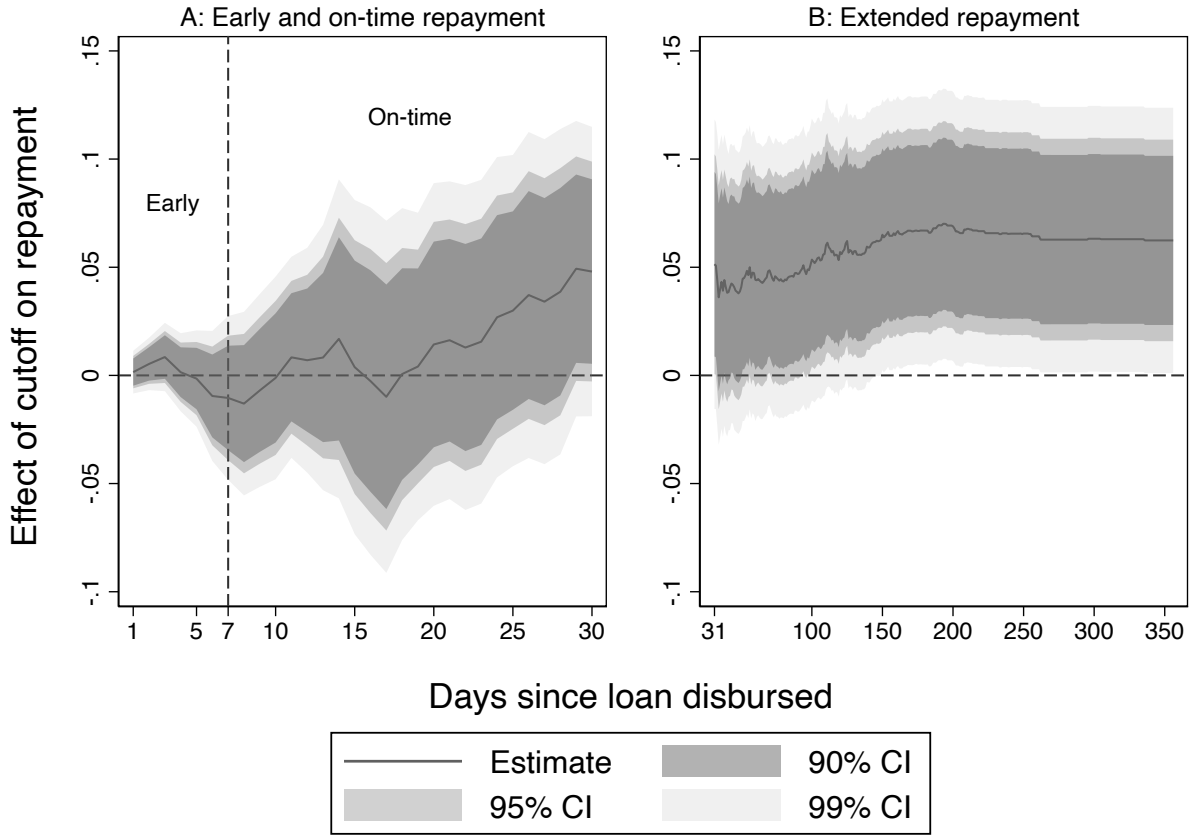
4.5 Timing of repayment

Our analysis thus far has considered the effect of delays on whether a loan is repaid. Now we study the timing of loan repayment. 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.²⁰ 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 at each post-disbursement date. By construction, the point estimate converges to the RD estimate in Table 1.

Figure 5 plots the RD estimates for each one of the 356 days. Panel A reports the estimates for 1-30 days after loan disbursement, and covers the normal 30-day repayment period. We label the period before seven days as the “early” repayment period because the

²⁰For example, a loan that is paid after 20 days is coded as follows: $Repaid = 0$ for days $t = 0, \dots, 19$, and $Repaid = 1$ for days $t = 20, \dots, 356$.

Figure 5: 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.

shortest possible loan term is seven days. Panel A clearly shows that there is no difference in the repayment behavior of delayed loans in the “early” repayment period. This indicates that clients are not returning their loans immediately after receiving them. Starting around 17 days post-disbursement, the slope of the estimate plot becomes positive, meaning that delayed loans are beginning to be repaid at a higher rate. After 30 days, we can detect a significant effect of the cutoff: 4.8 percentage points ($p = 0.064$), which means that roughly three-quarters of the overall delay effect has accrued by 30 days after disbursement.

Panel B reports the estimates for the repayment period 30-356 days after loan disbursement. During this period, the slope remains positive, explaining the remaining delay ef-

fect. While this confirms that for essentially any fixed window of 30 days or longer, delays have a significant positive effect on repayment, as a final check, we separately analyze loans submitted between November and January from later loans to ensure that unobservable loan repayments after the end of our observation period (which would be higher for later loans in our sample) could not counteract our results. Estimates are similar in both samples, suggesting that additional time to observe loan repayments would not change our findings.

4.6 Lender profits

While the welfare effects of loan delays are ambiguous for borrowers (see Section 6), our data suggest that delays can be used profitably by the lender. Consider the lender’s profit function, $\pi(\cdot)$, for a single loan:

$$\pi = p_r \cdot (L \cdot (1 + r) + F) - L - C \quad , \quad (2)$$

where p_r is the repayment rate, L is the loan amount, r is the interest rate, F is any fixed fee for the loan, and C is the lender’s marginal cost of the loan. If p_r increases by 8% –our main estimate– then revenue increases by 8%, and profit increases by more. Given that our estimates come from the set of already-approved loans, there is no ambiguity that the lender profits at the loan level from the delayed loans in our sample are higher.

Whether the lender is also better off at the borrower level over time depends on whether a delay affects the likelihood a borrower continues to take loans from the lender. If the answer is no, then the profit impact of a delay remains unambiguously positive. However, a delay affects the likelihood of future lending positively through an eligibility channel and negatively through a demand channel. The willingness of a lender to provide future loans depends on successful repayment of past loans. As delays increase the probability of repayment, they increase the set of continuing borrowers. On the other hand, the widespread advertisement of delivery speed in the digital credit industry suggests that borrower demand

will respond negatively to delays. Thus, the overall effect of a delay on lender profits over time is ambiguous.

In Appendix C, we show that there are *positive* but statistically insignificant effects of delays on the likelihood of borrowing again (4.8 percentage points with $p = 0.204$ for first loans, and 4.4 percentage points with $p = 0.391$ for second loans). When we condition the sample on successful repayment, the estimated delay effects are negligible (0.1 percentage points with $p = 0.982$ for first loans, and -0.8 percentage points with $p = 0.852$ for second loans). While these conditional demand estimates cannot be interpreted causally, they are not consistent with large negative effects of delays on future demand. Appendix Figure C1 shows estimated delay effects for a variety of other downstream loan demand outcomes. None of the estimates are statistically significant, and if anything, delays appear to increase future demand through higher repayment. Combined with the mechanically positive effects of delays on subsequent demand through the repayment requirement, this strongly suggests that in our sample the overall impact of delays on lender profit is positive.

5 Mechanisms

Several mechanisms might explain why disbursement delays increase digital loan repayments. In this section we evaluate a variety of potential options. Using the terminology of Karlan and Zinman (2009), the potential channels can be broadly broken down into hidden information and hidden action mechanisms. In the former, an extra delay induces a selection out of borrowing that improves the lender’s risk within the remaining pool. In the latter, an extra delay causes otherwise similar borrowers to make different choices. Within the set of hidden actions, we identify four different channels: the liquidity, opportunity, and behavior channels, and an additional channel related to the moral hazard of strategic default.

We argue that selection is unlikely to be at play in our setting. In addition, while all the hidden action channels might play a role, we find direct support in our data for the behavior

channel. We discuss each of these potential mechanisms below.

5.1 Hidden information

Selection could explain our main results if borrowers with low propensities to repay are more likely to cancel their loan requests when their loan deliveries take too long. To evaluate this mechanism, we obtained from the lender a separate dataset of successful applications that ended in the client rejecting their loan prior to disbursement. For the study period, we identified a total of 580 approved loans that were rejected by applicants prior to disbursement. These make up 2.5% of all loans disbursed. While we don't know when these withdrawn loans would have been disbursed, we know when their approved applications were submitted. We add these withdrawn loans submissions to our main database of issued loans and estimate the cutoff effect on an indicator equal to one if the loan was withdrawn.²¹ With a point estimate of -0.007 ($p = 0.363$), we rule out the possibility that delayed loans were more likely to be withdrawn.²²

While this result may be somewhat surprising, it is important to note that, once the loan has been fully processed and is awaiting disbursement by the lender, the borrower must contact the lender's customer service in order for the transaction to be canceled. Doing so takes some effort on the part of the client, and can explain why so few clients decide to withdraw the loan request.²³

²¹As the dataset of withdrawn loans does not contain information on borrower characteristics, we run the model on application time fixed effects only, corresponding to the specification in column (3) of Table 1.

²²This finding does not rule out the possibility that non-random selection might occur at other times in the borrowing process. For example, clients may decide to withdraw during the verification process (perhaps by refusing to answer the phone). Such loan declines do not affect our identification strategy, as they would occur prior to the loans being sent to batches. Furthermore, these calls are only for first-time borrowers.

²³Another way of "selecting out" could be clients returning delayed loans immediately after disbursement rather than holding them for their option value. However, only 8% of all loans were returned within seven days of disbursement, and Figure 5 shows that there is no difference in repayments between delayed and immediate loans disbursed in that time period.

5.2 Hidden action

Having ruled out selection, we hypothesize that the effect of a delay on repayment is the result of a change in the actions of borrowers. We discuss four hidden action channels below, starting with the behavior channel, which we argue fits the data best.

The behavior channel In contrast to delays changing the set of spending options available to the borrower, there are a variety of reasons why delays may affect borrowers' preferred spending options from a fixed choice set. For example, delays provide borrowers with extra time to deliberate about the use of their approved loans.²⁴ Existing research suggests that waiting periods can affect the consumption choices individuals make (Imas et al., 2016; DeJarnette, 2018; Brownback et al., 2019; Thakral and Tô, 2020), and that they do so in the direction of reducing “impulsive” or “tempting” choices.²⁵ In our context, increased deliberation could convince borrowers to change the use of the loan in a way that increases its expected return, leading to more liquidity at the time of repayment.²⁶

Alternatively, we speculate that without a delay, an individual may be able to apply for, obtain, and use a loan without confronting their partner, while household bargaining becomes an issue if disbursement is delayed, particularly overnight. As with individual deliberation, shifting the loan spending decision from individual to (at least in some part) joint may change the preferred loan use even if the set of available options doesn't change. There is a rich literature in labor, public, and development economics documenting the discrepancy between individual and household decision-making (Browning et al., 1994; Lundberg et al., 1997; Browning and Chiappori, 1999; Duflo, 2003; Bobonis, 2009; Attanasio and Lechene, 2014; Attanasio et al., 2012; Wang, 2014). A closely related thread of work demonstrates

²⁴This may be especially true if the extra delay is unexpected, as we believe is the case here.

²⁵Formal models of prospection (Gabaix and Laibson, 2022) and anticipation (Thakral and Tô, 2020) can both explain this effect.

²⁶This situation presumes 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).

a revealed preference for individual control over resources within households in developing countries, implying that those resources would be used differently in individual vs. joint decision-making (Anderson and Baland, 2002; Baland et al., 2011; Schaner, 2015; Jakiela and Ozier, 2016; Schaner, 2017). Having to explain and discuss the purpose of a loan with one’s partner could result in it being used differently. Especially if the original intended purpose was for a temptation good or negative expected return investment, we expect that joint decision-making would divert the loan into a higher expected return use.²⁷


Household bargaining has some clear testable hypotheses within the scope of our limited data; married borrowers that encounter overnight delays should be most affected by the delay. This is exactly what we find. Figure 6 shows estimates of the effect of missing a batch cutoff for a variety of different sub-samples. In particular, we report the effects by whether missing a cutoff induces an overnight vs. same-day delay, by marital status, and their interactions.²⁸ These estimates are reported in the Single Interactions section of the figure (the sample variable interacted with *PostBatch*). The cutoff effect is large and positive for overnight delays ($p = 0.011$), but zero for same day delays (difference: $p = 0.118$), and it is large and positive for married borrowers ($p = 0.003$), but near zero for unmarried borrowers (difference: $p = 0.047$). When we consider the interactions between the two measures (Double Interactions), we see that an overnight delay increases repayment for married borrowers ($p = 0.003$), whereas a same day delay has a negative, but imprecise effect ($p = 0.374$). We can reject that the overnight and same day coefficients for married borrowers

²⁷It could also be that when a loan is used for an agreed-upon expense, a partner with their own resources is more willing to help with repayment.

²⁸Partitioning the sample by whether the cutoff induces an overnight delay (or not) is akin to controlling for a post-treatment variable, which may cause bias. In this case we would expect a downward bias for the effect of overnight delays and an upward bias for the effect of same day delays. Alternatively, we can partition based on whether an application is submitted before or after noon, which has a big impact on the likelihood that missing a cutoff means an overnight delay. Results are similar in direction but smaller in magnitude: missing the cutoff on a loan submitted after noon has roughly three times the positive impact on repayment as missing the cutoff on a loan submitted before noon. Considering the sign of the bias and that overnight delays are the direct object of interest, we use that variable here.

are the same ($p = 0.014$). Given a potentially important role of household bargaining in explaining our results, we explicitly study differences in the married, overnight cutoff effects by gender (Triple Interactions). We find that overnight delays matter for both married men ($p = 0.078$) and women ($p = 0.003$), but that same day delays do not matter ($p = 0.247$ for married men, and $p = 0.908$ for married women). For both married men and women, we can reject that the overnight and same day coefficients are the same (male difference: $p = 0.048$, female difference: $p = 0.078$).

Figure 6: **Heterogeneous Treatment Effects of Interest for Household Bargaining**



Notes: Dashed vertical line represents our main estimate from column (4) of Table 1, and all estimates reported here use that same specification applied to the indicated sub-sample. While we do not estimate all these coefficients simultaneously with interaction terms, we use single, double, and triple interactions to refer to how many modifiers are used to define the sample.

The opportunity channel Borrowers facing time-sensitive consumption or investment opportunities that expire before loans are delivered might not want or need their loan by the

time it is received. Thus, higher repayments could be explained by the fact that funds have been unused (or used differently). This channel relies on a controversial assumption: it must be the case that on average, the financial return from loans that arrive quickly is lower than the return from loans that arrive with an extra delay. There are a variety of reasons why this may be the case. Time-sensitive investment opportunities might be particularly likely to be for high-risk low-reward investments (e.g. gambling). Time-sensitive loan demand might be more likely to be consumption-based than investment-based, stemming from temptation goods or emergencies. Without observing how borrowers use their loans, we can't evaluate this assumption directly. Rather, we note that *if the assumption holds*, extra delays can increase borrowers' ability to repay the loan by changing the available set of spending options when the loan arrives.

The heterogeneity analysis in Appendix Table A7 provides indirect support for the opportunity channel. Borrowers with more income and borrowers with better credit scores are more likely to have access to cheaper alternative sources of credit, but these sources would be slower than a digital loan. When such a borrower takes out a fast but high-cost digital loan, it is particularly likely that their intended purpose is time-sensitive, and thus how the loan is used will be particularly sensitive to the delay. Conversely, lower income and less creditworthy borrowers may lack access to slower, cheaper credit, and they may rely on digital loans for activities that are not time sensitive. When splitting the sample by above/below median income, we find a cutoff effect on repayment 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 (difference: $p = 0.275$). We also estimate 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" (difference: $p = 0.048$).

We highlight one particular time-sensitive, negative-expected-return investment opportunity that other work has identified as very relevant for digital credit and mobile money

(albeit in Kenya rather than Mexico). It is common to use digital credit and mobile money for gambling —sports betting in particular; from April to September 2021 the mobile money platform M-Pesa handled about 737 million USD in sports bets (Rosen, 2022). Chamboko and Guvuriro (2021) finds that a third of a representative sample of Kenyan digital borrowers were digital bettors, and that the bettors were more likely to have defaulted. While we do not have similar data from Mexico, we note that in 2019 there was approximately two billion dollars of revenue generated from regulated online gambling in Mexico, which is estimated to represent only 20% of the online gambling market (López, 2021).

The extra liquidity channel At the time of the study there were many other digital lenders offering similar products to the one of our lender. Thus, delayed borrowers with urgent needs could seek alternative sources of credit from other lenders. This additional credit could provide the necessary liquidity to repay delayed loans but at the cost of a higher level of overall debt. We cannot directly explore credit use with our data, as we only have administrative data from one digital lender. We note however, that are estimated cut off effects are concentrated among borrowers with high credit scores whose ability to borrow from other credit sources is highest. Thus, we do not rule out this channel.

Strategic default Carlson (2017) shows that some borrowers of a Kenyan digital lender engage in a type of strategic behavior: they move through loans quickly in an attempt to get to larger loan sizes before defaulting. This could interact with disbursement delay length in a variety of ways. If a loan takes longer to arrive than such a strategic borrower expects, this strategy may become less attractive. Alternatively, if a loan takes longer to arrive than a non-strategic borrower expects, they may decide that strategic default is the best use of credit from a digital lender. Another layer of ambiguity about the predictions of a strategic default model is that we don't know whether any given loan in our data is the loan a strategic defaulter plans to default on or is an intermediate loan they plan to repay. By focusing on first loans, the predictions are a little clearer. If first-loan delays shift borrowers out of the

strategic default build-up strategy, extra delays should reduce repayment. This is not what we find. If first-loan delays shift borrowers into the strategic default build-up strategy, extra delays should increase repayment, but by a smaller margin than for repeat loans. This is not what we find either.

An alternative assessment of strategic default comes from identifying borrowers who are more likely to be strategic types rather than trying to extract specific predictions from a model that suggests delays could generate both shifts into and out of a strategic default strategy. For example, perhaps delays change when strategic defaulters decide to default. Carlson (2017) shows that the median number of days between borrowers' first and second loans in her sample is one day, and she hypothesized that the speed at which a borrower takes out loans is correlated with strategic borrowing behavior. Thus, she splits the sample into two groups: borrowers who take out their second loan within one day of repaying their last loan, and borrowers who do so after a day or more. Indeed, among borrowers who take out loans quickly, getting an upgraded amount on the second loan relative to the first causes default rates to fall by around 100%.²⁹

In our own data, we also find that the median number of days between borrowers' first and second, or second and third loans is one day. We label the one day or less group "fast borrowers" and the remaining sample "slow borrowers." For fast borrowers, the cutoff effect (on the loan that is taken quickly) increases repayment by 4.4 percentage points, but the effect is not statistically significant. For slow borrowers, the delay increases repayment by 10.6 percentage points and the effect is statistically significant at the 1% level. Thus, we do not find that the effect of the delay on repayment is concentrated among borrowers we think are more likely to be strategic.

²⁹This suggests they intend to grow the loan even larger before defaulting.

6 Conclusion

We study whether one of the distinguishing features of digital credit—the rapid speed of delivery of funds—affects the likelihood that a loan is repaid. Despite the rapid growth of this market, this question remains unanswered. We combine difficult-to-obtain administrative data from a digital lender with a quasi-experimental identification strategy, and shows that reducing the speed of delivery of digital loans increases the likelihood that loans are repaid. Our preferred full-sample estimate indicates that roughly doubling loan provision speed from ten to 20 hours increased repayment by 6 percentage points. This effect corresponds to a 21% reduction in the likelihood of default.

This finding naturally raises the question of whether regulating the speed of digital credit disbursement, such as by imposing a waiting period on loan delivery, would improve consumer welfare. While the results in our setting show that such regulation could protect consumers from avoidable defaults, the full answer requires a careful welfare analysis. With the information at hand, this analysis is ambiguous. Several factors point to a positive impact on consumers. First, by avoiding defaults, clients’ credit scores are not hurt. Second, clients who repay maintain eligibility for future loans with the lender. This benefit appears important, as the demand for future loans from delayed borrowers is similar to the demand of borrowers who were not delayed. Finally, a number of the plausible mechanisms we discussed in Section 5 raise the possibility that borrowers may regret some of the most time-sensitive loan expenditures. On the other hand, we cannot rule out the possibility that some delays caused borrowers to fail to address immediate needs, led them to take additional loans from alternate sources that ultimately increased their overall indebtedness, or prevented them from taking advantage of timely and profitable opportunities. That said, the average wait for borrowers who just miss batch cutoffs is less than 24 hours, and for repeat borrowers it is about 11 hours. Thus, “delayed” digital credit is still extremely fast credit. Only very time-sensitive loan uses would imply welfare losses due to our observed delays in loan delivery.

We can be more conclusive about the effect of delays to the lender: profits are higher on delayed loans, and there is no evidence that they reduce future loan demand in our data. That said, based on discussions with both our lender partner and others in the industry it seems very likely that digital borrowers care about loan delivery speed to some degree. Thus, our broad takeaway from this study is that *targeted* waiting periods could be part of a lender's toolkit for broadening access to credit, rehabilitating past delinquent borrowers, or screening new borrowers based on their willingness to wait.

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.
- Alcott, H., J. Kim, D. Taubinsky, and J. Zinman (2021). Are high-interest loans predatory? theory and evidence from payday lending. *Review of Economics Studies*.
- Anderson, S. and J.-M. Baland (2002). The economics of roscas and intrahousehold resource allocation. *Quarterly Journal of Economics* 117(3), 963–995.
- 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.
- Attanasio, O., E. Battistin, and A. Mesnard (2012). Food and cash transfers: Evidence from colombia. *Economic Journal* 122(559), 92–124.
- Attanasio, O. and V. Lechene (2014). Efficient responses to targeted cash transfers. *Journal of Political Economy* 122(1), 172–222. Forthcoming, *Journal of Political Economy*.
- Baland, J.-M., C. Guirkinger, and C. Mali (2011). Pretending to be poor: Borrowing to escape forced solidarity in cameroon. *Economic Development and Cultural Change* 60(1), 1–16.
- 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.
- Bertrand, M. and A. Morse (2011). Information disclosure, cognitive biases, and payday borrowing. *The Journal of Finance* 66(6), 1865–1893.

- 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., J. E. Blumenstock, O. Folajimi-Senjobi, J. Mauro, and S. R. Nair (2022). Instant loans can lift subjective well-being:a randomized evaluation of digital credit in nigeria. *Working paper*.
- 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.
- Bobonis, G. J. (2009). Is the allocation of resources within the household efficient? new evidence from a randomized experiment. *Journal of Political Economy* 117(3), 453–503.
- Brailovskaya, V., P. Dupas, and J. Robinson (2021). Is digital credit filling a hole or digging a hole? evidence from malawi. Technical report, National Bureau of Economic Research.
- 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.
- Browning, M., F. Bourguignon, P.-A. Chiappori, and Lechene (1994). Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political Economy* 102(6), 1067–1096.
- Browning, M. and P.-A. Chiappori (1999). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica* 66(6), 1241–1278.
- 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.
- Chamboko, R. and S. Guvuriro (2021). The role of betting on digital credit repayment, coping mechanisms and welfare outcomes: Evidence from kenya. *International Journal of Financial Studies* 9(10), 1–12.
- 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. (2003). Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in south africa. *World Bank Economic Review* 17(1), 1–25.
- 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*.
- Gabaix, X. and D. Laibson (2022). Myopia and discounting. *Working 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.
- Jakiela, P. and O. Ozier (2016). Does africa need a rotten kin theorem? experimental evidence from village economies. *The Review of Economic Studies* 83(1), 231–268.
- 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.
- López, A. M. (2021). Key figures on the online gambling market in mexico as of 2019. *Statista September 15*.
- Lundberg, S. J., R. A. Pollak, and T. J. Wales (1997). Do husbands and wives pool their resources? evidence from the united kingdom child benefit. *Journal of Human Resources* 32(3), 463–480.

- 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.
- Rosen, J. W. (2022). How mobile money supercharged kenya’s sports betting addiction. *MIT Technology Review April 14*.
- Schaner, S. (2015). Do opposites detract? intrahousehold preference heterogeneity and inefficient strategic savings. *American Economic Journal: Applied Economics* 7(2), 135–174.
- Schaner, S. (2017). The cost of convenience? transaction costs, bargaining power, and savings account use in kenya. *Journal of Human Resources* 52(4), 919–945.

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*.

Totolo, E. (2018). Kenya’s digital credit revolution five years on. *CGAP BLOG March 15*.

Wang, S.-Y. (2014). Property rights and intra-household bargaining. *Journal of Development Economics* 107, 192–201. Working Paper.

A Appendix for online publication

Figure A1: Distribution of loan application submissions and loan disbursements






Figure A2: 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 A3: 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.

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 (0.003)	-0.004 (0.003)	-0.008 (0.005)	-0.009 (0.005)
	$p = 0.274$	$p = 0.218$	$p = 0.103$	$p = 0.081$
Age ²	0.000052 (0.000041)	0.000055 (0.000040)	0.000096 (0.000062)	0.000103 (0.000062)
	$p = 0.211$	$p = 0.172$	$p = 0.123$	$p = 0.099$
Female	0.013 (0.009)	0.011 (0.008)	0.018 (0.013)	0.018 (0.013)
	$p = 0.151$	$p = 0.177$	$p = 0.154$	$p = 0.154$
Married	-0.010 (0.010)	-0.011 (0.010)	-0.014 (0.015)	-0.015 (0.015)
	$p = 0.292$	$p = 0.251$	$p = 0.339$	$p = 0.318$
Dependents	-0.007 (0.004)	-0.006 (0.004)	-0.002 (0.007)	-0.001 (0.007)
	$p = 0.116$	$p = 0.142$	$p = 0.812$	$p = 0.937$
Log monthly income (pesos)	0.023 (0.005)	0.020 (0.005)	0.013 (0.008)	0.011 (0.008)
	$p < 0.001$	$p < 0.001$	$p = 0.117$	$p = 0.172$
Credit score (0-4)	0.026 (0.004)	0.034 (0.004)	0.065 (0.009)	0.073 (0.009)
	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
Log amount received (pesos)			0.029 (0.043)	0.003 (0.045)
			$p = 0.502$	$p = 0.955$
Loan term (days)			-0.003 (0.001)	-0.003 (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)				
<i>PostBatch</i>	6.56 (0.87)	10.76 (1.50)	9.85 (1.25)	9.81 (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)				
<i>PostBatch</i>	8.50 (1.57)	12.25 (2.00)	11.18 (1.77)	10.91 (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)				
<i>PostBatch</i>	6.71 (0.89)	8.60 (1.27)	8.34 (1.09)	8.25 (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)				
<i>PostBatch</i>	-0.148 (0.022)	-0.212 (0.043)	-0.231 (0.034)	-0.237 (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)				
<i>PostBatch</i>	-0.191 (0.037)	-0.284 (0.049)	-0.264 (0.040)	-0.261 (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)				
<i>PostBatch</i>	-0.161 (0.025)	-0.179 (0.035)	-0.201 (0.025)	-0.198 (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 (0.0013)	0.0043 (0.0018)	0.0042 (0.0017)
Estimate p -value	0.041	0.014	0.016
B. First-time loans (N = 3,090)			
Loan delay (hours)	0.0026 (0.0019)	0.0035 (0.0027)	0.0042 (0.0027)
Estimate p -value	0.172	0.193	0.123
C. Repeat loans (N = 4,087)			
Loan delay (hours)	0.0025 (0.0017)	0.0048 (0.0022)	0.0046 (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: OLS estimates of the impact of loan delay on loan repayment

	(1)	(2)	(3)
A. Full sample (N = 14,951)			
Loan delay (hours)	-0.0011 (0.0002)	-0.0012 (0.0002)	-0.0004 (0.0002)
Estimate p -value	0.000	0.000	0.080
B. First-time loans (N = 7,148)			
Loan delay (hours)	-0.0006 (0.0003)	-0.0006 (0.0003)	-0.0004 (0.0003)
Estimate p -value	0.028	0.037	0.211
C. Repeat loans (N = 7,803)			
Loan delay (hours)	-0.0003 (0.0003)	-0.0002 (0.0004)	-0.0002 (0.0004)
Estimate p -value	0.244	0.613	0.568
Day-of-week, hour-of-day, month FEs	N	Y	Y
Borrower controls	N	N	Y

Notes: Heteroskedasticity-robust standard errors are shown in parentheses below the estimates. The fixed effects include the hour-of-day, day-of-week, and month of application submission. In Panels A and C, column (3), 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 A7: **Heterogeneity in repayment effects by borrower finances**

	(1)	(2)
A: Income	Below median	Above median
<i>PostBatch</i>	0.028 (0.035)	0.079 (0.031)
Estimate <i>p</i> -value	0.424	0.011
Bias-corrected estimate <i>p</i> -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
B: Credit Score	None/Marginal/Average	Better/Best
<i>PostBatch</i>	0.033 (0.028)	0.142 (0.044)
Estimate <i>p</i> -value	0.251	0.001
Bias-corrected estimate <i>p</i> -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.

B Data construction

B.1 Batching identification

We do not observe batch 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. We exclude from the data loans that are processed in isolation of others and appear “unbatched” (not belonging to any particular batch). 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 B1). This drops 259 loans. Alternative approaches, including using k-means clustering to identify batched versus unbatched loans using the minimum distance to another loan, produce similar results.

Figure B1: **Time between loan deliveries in a batch**



Among the “batched” loans, we use a 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


that each day consists of six batches, and let the algorithm assign each loan to one of the six batches under that assumption. 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 (this initialization does much worse on the whole 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 B2 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).

To explore the robustness of our results to our batching procedures, here we show the results of an alternative that allows us to simulate a distribution of our main estimates. We switch from the deterministic initialization procedure described above to an entirely random initial batching with no post-processing of the batching results. For each initial batching (day X # of batches), we randomly initialize the algorithm 1000 times to create 1000 copies

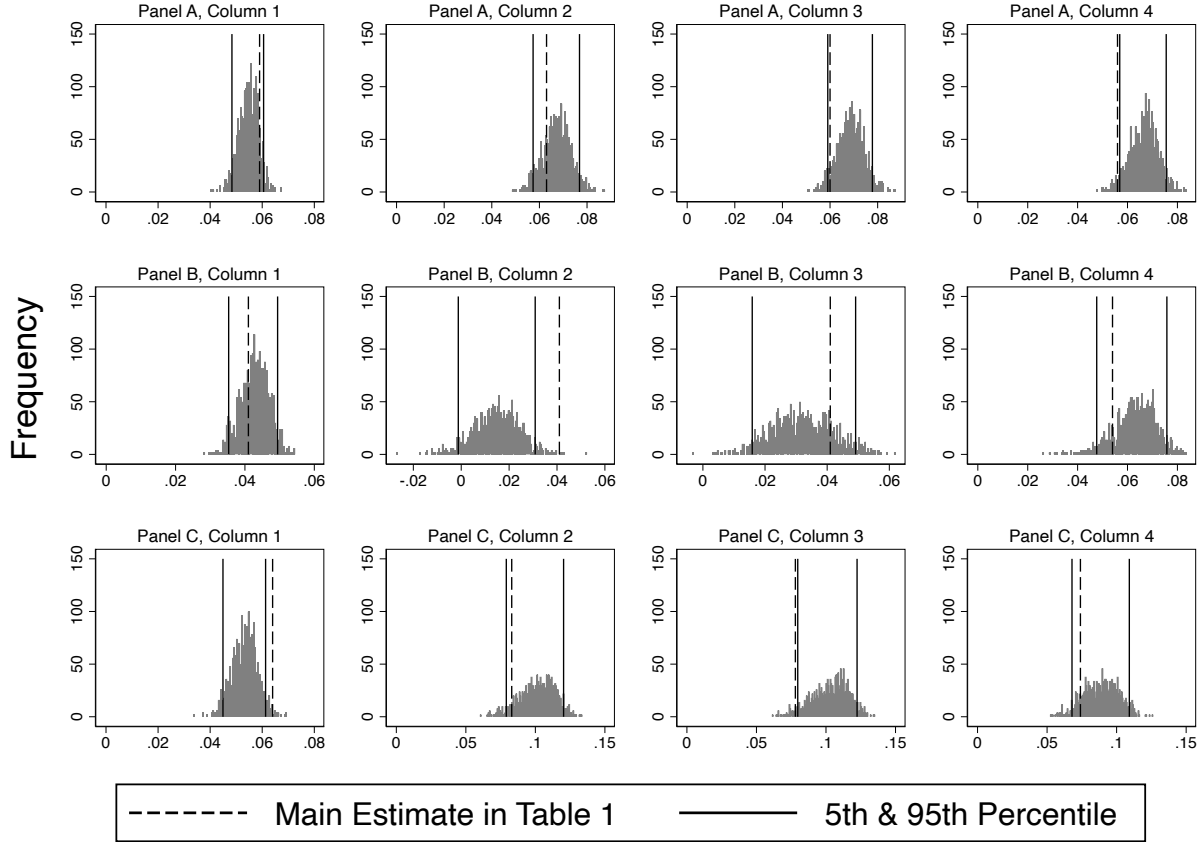
Figure B2: **Example of batching process**
December 14, 2018



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.

of our analysis sample. We then calculate distributions of estimates corresponding to each of the specifications reported in Table 1. Results are shown in Appendix Figure B3.

Figure B3: Estimate Densities from 1000 Random Initial Partition Batchings




Notes: Each graph corresponds to an estimate in Table 1. We randomly partition the submissions each day to initialize the batching algorithm, and skip any post-processing of batching errors. Then we follow identical procedures to obtain a regression-discontinuity estimate of the effect of missing the cutoff on the likelihood of repayment. We do this 1000 times in order to create these distributions. Estimates are reproducible using seeds 1-1000 in Stata 16.

B.2 Constructing the cutoffs

The lower-bound cutoff (LBC) is defined as the latest application submission time within a batch. Appendix Figure B4 shows the distribution of the LBCs in our sample.

We code *DistanceToBatch* as the difference (in minutes) between a loan’s application submission time and the nearest cutoff, and *PostBatch* as an indicator for whether *DistanceToBatch* is positive. Accordingly, in Figure 1, the application submission (start of the verification process) of all the loans is closer to the Batch A LBC than to the Batch B LBC. Therefore, the loans are assigned to the Batch A cutoff.

Figure B4: **Distribution of lower-bound cutoffs**

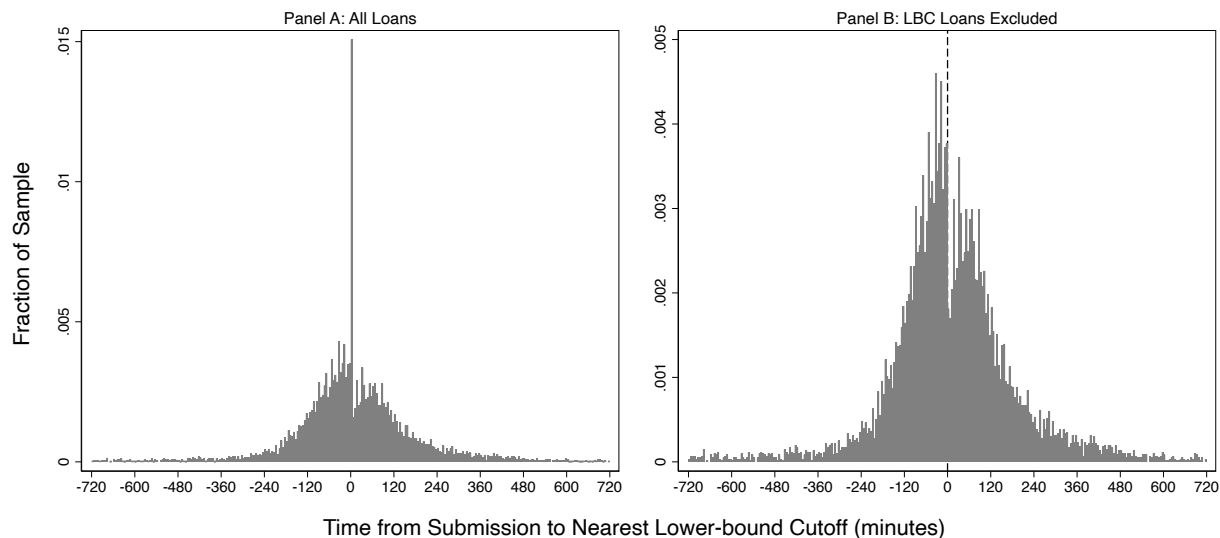


B.3 Selection around the cutoffs

Because the density of submissions drops after the LBC, the typical regression discontinuity validity test—smoothness of the density of the running variable—is not informative in our context. Appendix Figure B5, Panel B shows the density of loan submissions within a 12-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 over-represented in the distribution, shown in Panel A.

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

Figure B5: **Density of $DistanceToBatch$, 12-hour window**



Notes: Five-minute bins.

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 B5 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 B1 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.

Table B1: **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.

Does a 20-minute exclusion window make sense using other approaches? We now test directly for smoothness in a key observable –creditworthiness– through post-LBC cutoffs. Using the `rdrobust` optimal bandwidth approach, in Appendix Table B2, 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.

Finally, in Appendix Table B3, 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

Table B2: **Borrower credit score smoothness through post-LBC cutoffs**

Credit score: Minutes post-LBC	None		Marginal		Average		Better		Best	
	Coef.	<i>p</i> -value	Coef.	<i>p</i> -value	Coef.	<i>p</i> -value	Coef.	<i>p</i> -value	Coef.	<i>p</i> -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.

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.

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.

As an alternative to studying these post-LBC selection concerns in the estimation sample, we take a simulation approach to assessing ex-ante, how much selection bias should we expect given the structure of our data? First, we estimate the relationship between processing time and loan repayment outside the two-hour window around the LBC. We find that there is a 75.8% chance of repayment with no processing time, and a decrease of 0.1pp for each

Table B3: **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.

additional hour of processing time. This negative relationship (-0.1 p.p per hour) captures the variation in repayment that is in part due to covariates of processing time. Next, we take all loans inside the two-hour window of the LBC and replace their actual binary paid indicator with their predicted value of paid (which is now a likelihood) from that processing time regression. Then, we randomly resolve all of those likelihoods into paid indicators 1,000 different times. And finally, we estimate our RD with a variety of post-LBC exclusion windows for each of the 1,000 simulated samples. We repeat this exercise five more times, each time scaling up the estimated negative correlational relationship between processing time and repayment. We scale up the -0.1 estimate by factors of 1.1, 1.5, 2, 5, and 10. This allows for a larger impact of the processing time covariates on repayment than our baseline approach identifies.

In Appendix Figure B6 we report two statistics of the RD estimate distribution for each post-LBC exclusion window: the median estimate and the 95th percentile point estimate. The median tells us about the typical selection bias we should expect at each exclusion

window, and the 95th percentile tells us about when the outcome of no selection bias enters the confidence interval. For the median estimate, we find that without any scale-up, the magnitude of selection bias is always under 1 percentage point in absolute value (and is negative) for any post-LBC exclusion window. For comparison, our full-sample main estimate is roughly positive 6 percentage points. Scaling up, the selection bias remains negligible until we hit 5 times the observed relationship, at which point even a 40-minute exclusion window would lead us to substantially underestimate the causal effect of processing time. An exclusion window of 15 minutes or more is sufficient to fail to reject selection bias in all simulations excluding the one where we inflate the processing time-repayment correlation by 10 times.

Figure B6: Simulated Selection Bias for Loans Arriving Just After the LBC




Notes: Panel A shows the median point estimate from estimating the cutoff effect in simulated data without a treatment effect, and should be interpreted as the most likely degree of selection bias we would encounter for the given exclusion window and degree of endogeneity. Panel B shows the 95th percentile of the same distribution of point estimates that the median in Panel A comes from.

The use of an exclusion window after the LBC naturally raises the questions of whether there should also be an exclusion window before the LBC. If the LBC is selected based on a fast processing time, are loans just before the LBC also selected based on speed? The answer is no. This is because the processing speed of these loans has no bearing on either their likelihood of appearing in the sample or on whether they should be assigned to the pre- or post-cutoff period: as shown in Figure 2 only about 70% of these loans are disbursed in the same batch as the LBC. We can empirically corroborate this statement in two ways. First, we simply trace out whether loan delay decreases smoothly or discontinuously as the LBC is approached from the left in our data. This is shown in Appendix Figure B7, Panel A, where we put loans submitted within one hour before the LBC into five minute bins, take the bin means, and fit a linear prediction. We also show the LBC loan mean (at zero on the horizontal axis).

Second, we perform a similar exercise using simulated data to demonstrate that the discontinuous decrease in delay at the LBC arises from the structure and construction of our data rather than from the relationship between borrower characteristics and delay prior to the LBC. We randomly generate 100,000 loan submissions across a 2000 day sample with 50 loans per day. Within each day, we assign loans to three batches –morning, afternoon, and evening– where the batch disbursement times are uniformly drawn from 8-10am, 12-5pm, and 7-8pm, respectively. We then assign a delay to every loan by randomly sampling from the distribution of real delays, construct LBCs by identifying the latest submission time within each batch, and assign each loan to either the pre- or post-period of the nearest cutoff. As with the corresponding exercise with the observed loans we put loans submitted within one hour before the LBC into five minute bins, take the bin means, and fit a linear prediction. We also show the LBC loan mean (at zero on the horizontal axis). Results are in Appendix Figure B7, Panel B. In both the observed and simulated data it is clear that there is no selection of loans based on processing speed prior to the LBCs.

Figure B7: Discontinuous Change in Delay Length at the LBC



Notes: The observation at zero on the horizontal axis represents all LBC loans ($DistanceToBatch = 0$). Linear fits are obtained from OLS regressions applied to all loans (real and simulated in Panels A and B respectively) approved within one hour prior to an LBC.

C Impact of delays on the lender

Delays may have important knock-on effects on borrowers' demand for future loans. There are two potentially countering effects of first-loan delays on subsequent demand. On one hand, the lender offers additional loans only to clients who have paid. Because delayed loans are repaid more often, a delay increases the size of the client base. Although, 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.³⁰ On the other, 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 demand.


We estimate the effect of being delayed on the first loan on the likelihood of borrowing again. Using our preferred specification from column (4) of Table 1 we separately estimate

³⁰We 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.

the effect of delays in the first and second on post-repayment loan outcomes. Results are shown in Appendix Figure C1. We estimate the effect of a first loan delay on all subsequent outcomes: taking a second loan, repaying a second loan, taking a third loan, and repaying a third loan. All of these estimates are positive, and the effect on repaying a second loan and taking a third loan are marginally statistically significant ($p = 0.099$ and $p = 0.069$ respectively). The effect of a second loan delay on taking a third loan and repaying a third loan are also both positive, but neither is precise.

We also estimate the effect of first and second loan delays, conditioning on intermediate outcomes. While this confounds causal identification by inducing endogenous control bias, we argue that the estimates are still informative for ruling out large negative demand effects of delays. For both first and second loan delays, we find that among borrowers who repaid, the delay has zero impact on the likelihood of borrowing again. Whatever negative demand effects exist are not large enough to show up in these specifications.

Figure C1: Impact of delays on post-repayment outcomes



Notes: All estimates reported here use the same RD specification as in Table 1, column (4).