Do Firms' Behavioral Biases Delay the Search for Credit During the Covid-19 Crisis?*

Pre-Analysis Plan

Paul Gertler Sean Higgins Ulrike Malmendier Waldo Ojeda

June 6, 2020

Abstract

The goal of this study is to estimate the effect of deadlines and reminders on firms seeking out credit during an economic downturn due to the Covid-19 outbreak. Credit can be valuable to firms to weather the downturn, but its benefits are in the future while the cost to apply for credit is borne immediately. Forgetfulness and present-bias might be biases that exacerbate the cost and lead firms to not ever search and request a loan. We will study if treatments involving deadlines and reminders (anticipated and unanticipated) help firms overcome these biases to search for a loan.

1 Introduction

In this document we provide additional details not included in the pre-registration submitted April 29, 2020 to the AEA RCT Registry (AEARCTR-0005782). The intervention, primary outcomes, experimental design, experiment characteristics (such as sample size and how the randomization was conducted), and IRB details were all submitted April 29, 2020, which was one day prior to the start of the RCT. Here we include those details as well as additional details uploaded prior to accessing data on experimental outcomes from our implementing partners.

*We thank the CEGA-Visa Financial Inclusion Lab for research funding and Miguel Angel Jimenez, Alexandra Wall, and Tiange Ye for research assistance. IRB approvals: UC Berkeley IRB 2018-02-10796 and 2020-03-13091.

2 Design

2.1 Intervention

In our intervention, we will contact firms to provide them details about the opportunity to apply for a loan from a FinTech lender. The firms in our sample are clients of a FinTech electronic payment processing company that is partnering with a separate FinTech lender. The firms will be informed about the loan opportunity through a concurrent email and SMS messages. To test if forgetfulness and present-bias hinder the search for credit, we will vary the content and timing of the emails and SMS messages: firms will be randomly assigned offers to apply for a loan with different combinations of deadlines and unanticipated and anticipated reminder messages.

2.2 Primary outcomes

Our primary outcome is a dummy variable indicating whether firms apply for a loan. We will also use this outcome to estimate the model in Ericson (2017) to quantify the relative importance and interaction of different behavioral biases in preventing firms from applying for credit.

2.3 Secondary outcomes

Through link tracking, we know the fraction of firms in each treatment group that click the links (but cannot link these back to particular firms within each treatment group). Thus, a secondary outcome will be clicking the link. In addition, for those who log in after clicking the link, the FinTech lender also tracks how much progress firms make in completing the loan application (and these measures we can link to other data on the firms from the FinTech payments provider): 0% complete if after clicking the link they log in but do not fill out general information, and 25%, 50%, or 75% complete if they do not fully complete the loan application but complete a fraction of the application. 100% complete applications are our primary outcome "apply for a loan" above. Dummy variables for each of the four partially complete outcomes (0%, 25%, 50%, 75% complete) will be used as secondary outcomes to explore firms' behavior in the loan application process.

2.4 Experimental design

Messages with deadlines will state that the firm has one week or one day to access the link to apply for credit. Messages with an unanticipated reminder will receive a reminder to access the link one week after. Messages with an anticipated reminder will also receive a reminder to access the link one week after and will be told in the initial email and SMS messages that they will receive this reminder and on what day they will receive it. In total, our design has eight treatment groups:

1. Control, no messages
2. Loan messages with no deadline, no reminder

3. Loan messages with no deadline, anticipated reminder

4. Loan messages with no deadline, unanticipated reminder

5. Loan messages with 1-week deadline, no reminder

6. Loan messages with 1-week deadline, anticipated reminder

7. Loan messages with 1-week deadline, unanticipated reminder

8. Loan messages with 24-hour deadline, no reminder

2.5 Randomization method

Randomization done by a computer (R script).

2.6 Randomization unit and sample size

We randomize at the individual firm level (note that outcomes are measured at the firm level) using a simple stratified randomization. There are 70,020 individual firms in our experiment.

2.7 Stratification

We stratify our randomization by four variables:

1. Average monthly electronic sales in the past year, or since the firm registered with the payments processing company if it was within the past year (4 quartiles)

2. Business type (6 categories: Beauty, Clothing, Professionals, Restaurants, Small Retailers and Other)

3. Tax registration status (2 categories: Self-Employed and Limited Company)

4. A proxy for initial impact on sales due to the Covid-19 outbreak as of March 2020. This proxy is defined as above or below the median in the percent difference in sales from February 2020 to March 2020. A third group for this categorical variable is made up of those who had no sales in February 2020 such that the change in sales measure is undefined.

We stratified on these variables since we will test for heterogeneous treatment effects by each of these variables (sales, business type, tax registration status, and how affected they were by Covid-19). Our stratification includes 144 blocks: 4 (prior sales quartiles) * 2 (tax registration status) * 6 (business types) * 3 (impact on sales from Covid-19 outbreak).

2.8 Experiment characteristics

Sample size by treatment arm.

1. Control, no messages: 327
2. Loan messages with no deadline, no reminder: 10,355
3. Loan messages with no deadline, anticipated reminder: 8,592
4. Loan messages with no deadline, unanticipated reminder: 10,362
5. Loan messages with 1-week deadline, no reminder: 10,765
6. Loan messages with 1-week deadline, anticipated reminder: 8,104
7. Loan messages with 1-week deadline, unanticipated reminder: 10,755
8. Loan messages with 24-hour deadline, no reminder: 10,760

Power calculations. Sample sizes for each group were informed by results in a prior pilot test conducted in May 2019 and making pairwise power calculations for comparisons of interest. In the pilot we were offering a reduction in the merchant fee charged to the firm for card payments processed through the FinTech payments company from 3.75% to 3.5%. There was a control group, a placebo group that received an email from the FinTech payments company with no messages, a group that received the offer with no deadline, and a group that received the offer with a 24-hour deadline. We also sent two non-randomized reminders to both groups after the deadline had passed (the deadline was not binding, so firms could still sign up after the deadline).

For the purpose of the power calculations we assume a similar take-up between a 24-hour and one week deadline absent other treatments. The take-up rates in the pilot by pairwise comparison and the necessary sample size to detect such differences are as follows (using the treatment group numbers above). In the comparisons below, P0 refers to take-up in the second group listed in each comparison (e.g. if the comparison is 2 vs 1, P0 refers to take-up in group 1, and P1 refers to take-up in group 2). For the reminders which were not randomized in our pilot, we estimate P1 as cumulative take-up of the offer 24 hours after the reminder was sent and P0 as cumulative take-up immediately before the reminder was sent.

- 2 vs 1: P0 = 0.01, P1 = 0.18. Minimum sample size per arm = 46.
- 5 vs 1: P0 = 0.01, P1 = 0.28. Minimum sample size per arm = 26.
- 5 vs 2: P0 = 0.18, P1 = 0.28. Minimum sample size per arm = 277.

- 4 vs 2: P0 = 0.14, P1 = 0.18. Minimum sample size per arm = 1472.
- 7 vs 5: P0 = 0.20, P1 = 0.24. Minimum sample size per arm = 1529.

In the pilot there was no anticipated reminder treatment group. To obtain an estimate of expected effect size of the anticipated reminder and perform power calculations, we benchmark results from the pilot with model simulations based on the model in Ericson (2017) assuming standard magnitudes for present-bias and forgetfulness from the literature, and assuming full naïveté about present-bias but accurate beliefs about memory (using Ericson's terminology, $\beta = 0.9, \hat{\beta} = 1, \rho = 0.95, \hat{\rho} = 0.95$). The model simulations predict ratios of the difference in take-up between the groups with unanticipated and anticipated reminders over the difference in take-up between the groups with unanticipated and no reminders. In our pilot we use the difference in take-up between groups with unanticipated and no reminders and then apply the ratio to scale an estimated take-up rate for groups with an anticipated reminder. The ratio from the model is measured in the period when the reminder is sent right before the deadline. With these simulated take-up rates of groups with an anticipated reminder, we get a treatment effect ratio of 1.23, which means the necessary sample size to detect differences in relevant pairwise comparisons are:

- 4 vs 3: P0 = 0.14, P1 = 0.18. Minimum sample size per arm = 1222.
- 7 vs 6: P0 = 0.24, P1 = 0.29. Minimum sample size per arm = 1152.

To obtain a minimum sample size per arm for our study, we select the largest sample size needed for each group depending on its relevant pairwise power calculations above. For group 8, we assume we need the same number of observations as for group 5, since both include a deadline and no reminder. We calculate the following minimum sample sizes per arm to detect the expected effect sizes based on our pilot and simulations:

- Control, no messages: 46
- Loan messages with no deadline, no reminder: 1,472
- Loan messages with no deadline, anticipated reminder: 1,529
- Loan messages with no deadline, unanticipated reminder: 1,222
- Loan messages with 1-week deadline, no reminder: 1,152
- Loan messages with 1-week deadline, anticipated reminder: 1,472
- Loan messages with 1-week deadline, unanticipated reminder: 1,529

- Loan messages with 24-hour deadline, no reminder: 1,529

Thus, we need in total 9,951 observations across all treatment arms to statistically detect the expected differences in take-up between treatment groups of interest, based on outcomes in our pilot and simulations of the Ericson (2017) model. As our available sample size for the experiment is 70,020 firms, which is much larger than the needed 9,951, we adjust sample sizes of each treatment arm proportionally to arrive at the sample sizes per arm in our study shown under "Sample size by treatment arm."

3 Analysis Plan

3.1 Reduced form

Our experimental design allows to compare in reduced form the effects of a reminder or deadline on the probability of applying for a loan. Our primary results will be from the following regression:

$$y_i = \lambda_{s(i)} + \beta T_i + \varepsilon_i \tag{1}$$

where y_i is the outcome of interest, $\lambda_{s(i)}$ are strata fixed effects for the 144 stratification blocks defined above (which also absorb the constant), T_i is a vector of indicator variables denoting treatment assignment, and ε_i are heteroskedasticity-robust standard errors (not clustered since the randomization unit is the individual firm).

Note that for one of our secondary outcomes, clicking the link, we are unable to merge with data other than treatment assignment since the link clicks are tracked separately; thus for that outcome we will estimate (1) without strata fixed effects and with a constant α in their place.

Our main estimates of the effect of a reminder, deadline, and anticipated vs. unanticipated reminder will be estimated as follows:

- To estimate the overall effect of a reminder, the comparison of interest is between groups 3, 4, 6, or 7 (groups that receive a reminder) and groups 2 or 5 (groups that also receive the offer and comparable deadlines, but no reminder). For maximum power, we will also include groups 1 and 8 in the regression and thus use a vector of dummies for T_i for (i) group 1, (ii) group 8, and (iii) reminder groups i.e. groups 3, 4, 6, or 7. The omitted dummy will be for no-reminder groups 2 or 5 and the coefficient of interest will be the coefficient on the reminder dummy.

- To estimate the overall effect of a deadline, the comparison of interest is between groups 5, 6, or 7 (groups that receive a one-week deadline) and groups 2, 3, or 4 (groups that also receive the offer but no deadline). For maximum power, we will also include groups 1 and 8 in the regression and thus use a vector of dummies for T_i for (i) group 1, (ii) group 8, and

(iii) deadline groups 5, 6, or 7. The omitted dummy will be for no-deadline groups 2, 3, or 4 and the coefficient of interest will be the coefficient on the deadline dummy.

- To estimate the effect of an anticipated vs. unanticipated reminder, the comparison of interest is between groups 3 and 6 (anticipated reminder) vs. groups 4 and 7 (unanticipated reminder). For maximum power, we will also include the other groups in the regression and thus use a vector of dummies for T_i for (i) group 1, (ii) comparable no-reminder groups (2 and 5); (iii) group 8; (iv) anticipated reminder groups (3 and 6). The omitted dummy will be for unanticipated reminder groups (4 and 7) and the coefficient of interest will be the coefficient on the anticipated reminder dummy.

For robustness, we will also estimate the above three regressions with the groups that are not part of the comparison of interest dropped, and with T_i as a single dummy equal to 1 for either a reminder, a deadline, or an anticipated reminder.

In addition to the above broad comparisons, to make pairwise comparisons across the 8 treatment arms we will estimate (1) including the full sample and with T_i as a vector of 7 dummies for treatment arms 2 through 8 (with group 1, the control, as the omitted dummy).

We will estimate the above (i) for applications aggregated over the 8-day period between the initial messages and the deadline; (ii) for overall applications from the beginning of the experiment to the latest date in which we receive data from iZettle (beyond one week); (iii) we will estimate daily effects over the eight-day experiment period by estimating separate regressions for each day where the outcome is a dummy for loan applications; and (iv) weekly effects by estimating separate regressions for each week since the beginning of the experiment. For items iii and iv we will use two specifications: one with *cumulative* applications by that day or week and another with applications just in that day or week as the outcome.

Finally, to view take-up over time across treatment arms we will graph cumulative distribution functions of take-up by treatment arm and by combinations of arms (where combinations of arms refer to, for example, “reminder” referring to any firm in arms 3, 4, 6 or 7). We will also estimate a proportional hazards model where the dependent variable is the hazard rate of our outcomes of interest.[1] Coefficients on the treatment assignment variables when estimating this model will reveal how each treatment varies the proportion of take-up over time.

3.2 Heterogeneity

We will estimate

$$y_i = \lambda_{s(i)} + \beta T_i + \gamma T_i \times H_i + \varepsilon_i \tag{2}$$

[1]The specification is: $\lambda(t|T_i) = \lambda_0(t)exp(\beta T_i)$ where $\lambda(t|T_i)$ is the hazard rate of our outcomes of interest, $\lambda_0(t)$ is the baseline hazard function and T_i is a vector of indicator variables denoting treatment assignment.

where H_i is a vector of dummy variables that measure heterogeneity and γ is the vector of coefficients of interest. There is no simple (non-interacted) H_i term in (2) because it is absorbed by the strata fixed effects (since each strata is an interaction of the variables we will use for heterogeneity tests). Our main heterogeneity tests will be the variables that we stratify on which are described in more detail above, namely:

- Baseline sales quartiles (4 categories). (For increased power, we will also use above/below-median sales in a separate regression.)
- Business type (6 categories).
- Tax registration status (2 categories).
- Initial impact of Covid-19 on sales (3 categories: above median impact, below median impact, or undefined impact for firms with 0 sales in the comparison month of February 2020).

We will estimate (2) for the primary outcome of loan applications as well as the secondary outcomes of partial applications (0%, 25%, 50%, and 75% complete), but not for clicking the link included in the email and SMS messages as we cannot merge the data on clicks with firm characteristics.

3.3 Structural model

We will follow Ericson (2017) in our estimation method to structurally estimate forgetfulness and present-bias parameters based on the experiment outcomes. In the model, an agent makes a decision to perform a task that is beneficial in the future but has an immediate cost.

The agent has present-biased preferences and possibly naïveté: $U = u_0 + \beta\left(\sum_{t=1}^{\infty} \delta^t u_t\right)$, where δ is the discount factor, β is the present bias parameter, and the agent has beliefs $\hat{\beta}$ about β. The model also incorporates imperfect memory. There is a probability of remembering the task in period t conditional on remembering it in period $t-1$, measured by the parameter ρ_t (with $\rho_0 = 1$). The agent will only be able to perform the task if they remember it. Agents have beliefs $\hat{\rho}_t$ and are overconfident about their memory if $\hat{\rho}_t > \rho_t$. Reminders about the task raise ρ_t in the period they are sent. However, only an anticipated reminder in period t (when the agent is told about a reminder that it will receive in a future period t) increases the agent's expectations in earlier periods of $\hat{\rho}_t$.

In each period, the agent draws a cost c_t from a known distribution of costs $F(c)$ and would receive benefit y next period if they complete the task. Thus, the agent decides to act based on the current value function

$$W_t = \begin{cases} \beta\delta y - c_t, & \text{if act,} \\ \hat{\rho}_{t+1}\beta\delta E_t[V_{t+1}], & \text{if do not act,} \end{cases}$$

where $E_t[V_{t+1}]$ is the perceived continuation value of not doing the task in the current period (and potentially doing the task in a future period). Note that present bias leads the current value function W_t, which is a function of β, to differ from the perceived continuation value V_t, which is a function of $\hat{\beta}$. At the deadline, the continuation value is zero as the opportunity to perform the task in future periods is removed. By backwards induction from the deadline, the model leads to a cutoff strategy where the agent adopts in period t if the cost draw c_t is below a threshold c_t^*.

Specifically, by backwards induction we obtain a recursive set of expressions that implicitly define the cost threshold:

$$\begin{aligned}
c_t^* &= \beta\delta\left(y - \hat{\rho}_{t+1}E_t\left[V_{t+1}\right]\right) \qquad (3)\\
E_{t-1}\left[V_t\right] &= F\left(\hat{c}_t^*\right)\left[\delta y - E\left[\hat{c}|\text{act}\right]\right] + \left(1 - F\left(\hat{c}_t^*\right)\right)\delta\hat{\rho}_{t+1}E_t\left[V_{t+1}\right]\\
E\left[\hat{c}|\text{act}\right] &= \int_0^{\hat{c}_t^*} c dF(c)\\
\hat{c}_t^* &= \hat{\beta}\delta\left(y - \hat{\rho}_{t+1}E_t\left[V_{t+1}\right]\right)
\end{aligned}$$

The probability of adopting at period t is:

$$Pr\left(\text{adopt at } t\right) = \underbrace{F\left(c_t^*\right)}_{Pr(\text{worthwhile at } t)} \underbrace{\prod_{j=1}^{t}\rho_j}_{Pr(\text{remember})} \underbrace{\prod_{k=0}^{t-1}\left(1 - F\left(c_k^*\right)\right)}_{Pr(\text{didn't already adopt})} \qquad (4)$$

Thus, integrating over individual firms (whose i subscript was excluded above for ease of notation)—which can have heterogeneous costs—provides a set of moments giving the fraction of individual firms that adopt in period t, with one moment for each of the T periods (where T is the period in which the deadline occurs).

In our experimental setting, the benefit for firms is a loan from a Fintech company during the COVID-19 crisis and the task is clicking the link in the email and applying for a loan. The time period is a day. We set our deadline to be in one week (5pm on day 8 of the experiment) for all treatment groups that have a deadline except group 8. If assigned, the anticipated or unanticipated reminder is sent at the beginning of the day of the deadline (day 8). For each treatment group, we will obtain take-up rates for eight periods that will be used as moments to estimate forgetfulness and present-bias parameters. To isolate variation in costs from the probability of forgetting, we add a treatment group with an immediate one-day deadline (group 8), so that this group only faces a one-period decision and (under the assumption that each period is a day) does not risk forgetting. Our experiment thus allows us to estimate a set of moment equations of the form (4) to estimate β, $\hat{\beta}$, ρ, $\hat{\rho}$, and heterogeneous costs, where each treatment arm provides T moments.